



RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide



Contact Information

RSA Link at <https://community.rsa.com> contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

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

License Agreement

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

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

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

Third-Party Licenses

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

Note on Encryption Technologies

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

Distribution

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

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

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

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

May 2021

Contents

Decoder and Log Decoder Quick Setup	11
Configure Common Settings on a Decoder	14
Configure Capture Settings	16
Select a Network Adapter	16
Configure a Decoder to Begin Capturing Data Automatically	18
Configure Optional Capture Settings	19
(Optional) Configure System-Level (BPF) Packet Filtering	21
Examples	23
Testing	23
Conversions	24
(Optional) Configure a Decoder to Capture Data Across All Types of Network Interfaces	24
(Optional) Configure Meta-Only Decoders	27
(Optional) Configure Selective Network Data Collection	28
Use Predefined Policies	29
Create New Policies from Predefined Ones	31
Create Custom Policies	32
Verify Published Policies	33
Troubleshooting Policy Deployments That Fail	35
Unpublish Policies	36
Delete Policies	36
Supported Protocols for Selective Network Data Collection	37
(Optional) Configure a Decoder to Write Standard pcap-formatted Files	39
(Optional) Multiple Adapter Packet Capture	40
Configure Multiple Adapter Packet Capture	40
Per Interface Configuration	41
Stat Nodes	41
Per-Interface Stats	41
Aggregate Stats	42
Pool Stats	42
capture.port Meta Key	42
Special Devices	43
Packet arrived out of order Message	43
Multi-Assembler Modes	43
(Optional) Data Plane Development Kit Packet Capture	44
Advantages	44
How it Works	44

Migrate PF_RING Devices to DPDK	45
Manually Move a Capture Interface to DPDK	46
Utilizing Receive Side Scaling with DPDK	47
Advantages	47
Configuring Receive Side Scaling	48
How it Works	48
(Optional) Preserve VLAN Tags When Using the Packet MMAP Capture Interface	48
(Optional) Process Raw Syslog Data without Priority Field	53
(Optional) Configure Decoder to Support OpenAppID	54
Enable and Disable Parsers and Log Parsers	57
Start and Stop Data Capture	60
Configure Decoder Rules	61
Rule Processing	61
Rule and Query Guidelines	62
Rule Examples	62
Invalid Rules	63
General Syntax Guidelines	63
Capture Rule Syntax	64
Configure Capture Rules	66
Import Rules from a File and Export Rules	68
Push Rules to Other Services	70
Change Execution Order of Rules	72
Restore a Rule Snapshot from History	72
Configure Application Rules	74
Configure Correlation Rules	78
Configure Network Rules	82
Supported Meta Keys in Network Rule Conditions	82
Fix Rules with Invalid Syntax	87
Decoder Commands for Managing Rules	89
add Command	89
merge Command	90
Methods of Sending a List of Rules to a Service	90
Send a NetWitness Rule File	91
Send Numbered Parameters	91
Ordering Rules When Pushing	92
replace Command	93
clear Command	93
delete Command	93
validate Command	93

Configure Parsers and Feeds	94
Configure Parsers	94
Upload and Delete Custom Parsers	95
Upload Parsers to a Decoder or Log Decoder	95
Manage Upload Jobs	97
Delete Deployed Parsers	97
Enable and Configure the Entropy Parser	98
Entropy Parser Configuration in the Concentrator Custom Index File	101
Flex Parsers	103
Arithmetic Functions	104
Language Definition	104
Common Parser Operations	106
Match Port and Identify Immediately	106
Match Port and Delay Identification	106
Match Token and Identify Immediately	107
Match Multiple Tokens	107
Match Token and Create Metadata	108
General Functions	109
General Functions Language Definition	109
Logging Functions	111
Language Definition	111
Nodes	112
Nodes Language Definition	112
Payload Functions	116
Language Definition	116
Regex	118
Language Definition	118
String Functions	119
String Functions Language Definition	119
GeoIP2 Parsers	122
Lua Parsers	124
HTTP Parsers	125
Visibility into HTTP/2 Sessions	126
Decoder Snort Detection	127
Using Snort Rules in Decoders	127
Initial Decoder Snort Configuration	127
Snort Parser Output	129
Rule Statistics	130
Configuration Directives	130
Variable Definitions	130
ruletype	131

Other Configuration Entries	131
config detection: debug	131
config nopcre	131
config classification	131
Snort Parser Capabilities	131
Rule Sections	131
Content	132
Flow	133
Fast Pattern	133
PCRE	133
HTTP Parsing	133
File Data	134
Binary (byte) Directives	134
Default option / fallback	135
Packet Scanning/Scoping	135
Known Limitations for Snort Parser	135
Decoder Token Match Limits	135
Flow Bits	135
DCE Preprocessor	135
Thresholds and Filters	135
Raw Byte Scanning	135
IPv6	136
Performance Considerations	136
Content Tokens	136
Negation	136
Search Parser	137
Search Methods	137
Syntax	137
Parameters	137
Example	138
Wireless LAN Configuration	139
Troubleshooting Parsers	139
Lua Parser Errors	139
Results of Automatically Disabling a Parser	140
Resetting the Parser	140
Configure Feeds	140
Custom Feed Definition File Structure	141
Sample Feed Definition File	141
Define Multiple Values in a Single Field	142
Feed Definition Equivalents for Custom Feed Wizard Parameters	142
Sample Files for a MetaCallback Feed Using CIDR Index Range for IPv4 and IPv6	144
Feed Definitions File	146
Create a Custom Feed	147

Create a STIX Custom Feed	158
Create an Identity Feed	174
Import the SSL Certificate	178
Cannot Verify Identity Feed URL	179
Edit, Upload, or Remove a Feed	180
Create Custom Meta Keys Using a Custom Feed	184
Add a Custom Meta Key in the Log Decoder	184
Deploy a Log Decoder Feed in Live	185
Add the Custom Meta Key Entry in the Concentrator Custom Index file	188
Investigate a Custom Meta Key	189
Additional Procedures	189
Verify the Custom Meta Keys on ESA	189
Update the Schema in Archiver	190
Update the Schema in Warehouse Connector	191
Update the Schema in Reporting Engine	191
Decoder and Log Decoder Additional Procedures	193
Configure 10G Capability (Version 11.5 and Earlier)	193
Hardware Prerequisites	194
Software Prerequisites	194
Install the 10G Decoder	195
Download and Update the BIOS	195
Locate the 10G Decoder Packages	195
Verify 10G Decoder Packages Are Installed	195
Configure the 10G Decoder	195
Typical Configuration Parameters	197
Storage Considerations	197
Using the Series 4S Hardware (With Two or More DAC Units)	198
Using SAN and Other Storage Configurations	198
Capturing 10G with 40G Network Fabric	198
Parsing and Content Considerations	198
Best Practices	198
Tested Live Content	199
Aggregation Adjustments Based on Tested Live Content	200
Optimize Read/Write Operations When Adding New Storage	200
Configure High Speed Packet Capture Capability (Version 11.6 and Later)	203
Hardware Prerequisites	204
Software Prerequisites	204
Decoder Installation	204
Install the Decoder Service on Host that Will Perform the Capture	205
Select the Capture Interfaces and Assign Interfaces to DPDK	205
Verify the Decoder Packages Installation	205

Select the Decoder Operating Mode	205
Configure the Decoder for Normal Mode (Default Mode)	206
Configure the Decoder for 10G Mode	206
Performance Tuning Parameters	207
Storage Considerations	208
Packet Retention Requires Extremely High Sustained Throughput	208
Using SAN and Other Storage Configurations	208
Optimize Read/Write Operations When Adding New Storage	208
Parsing and Content Considerations	210
Best Practices	210
Use Case 1: 10G Mode - Egress, Deep Packet Inspection	211
Tested Live Content	211
Not Tested	212
Other	212
Aggregation Adjustments Based on Tested Live Content	212
Configure the Decoder for NDR Mode	212
Performance Tuning Parameters	214
Storage Considerations	215
NDR Mode Assumes No Packet Retention	215
Parsing and Content Considerations	215
Parsing at Speeds Greater than 10G	215
Best Practices	215
Use Case 1: NDR Mode - Egress, General Purpose	216
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + All Native Parsers	216
Tested Live Content	216
Not Tested	216
Other	217
Use Case 2: NDR Mode - Egress, Data Exfiltration	217
Generate NetFlow Style Meta Only + Exfiltration Focused Specific Native Parsers	217
Tested Live Content	217
Not Tested	217
Other	217
Use Case 3: NDR Mode - Lateral Movement	217
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + Specific Native Parsers	217
Tested Live Content	218
Not Tested	218
Other	218
Configure a Log Decoder to Accept Protobuf	219
Configure Session Split Timeouts	221
Configure Syslog Forwarding to Destination	224
Configure Transaction Handling on a Decoder	226
Transactions Off	226

Transactions Represented as Meta Items	227
Transactions Split Sessions	227
Configure Data Export	228
Decrypt Incoming Packets	229
Step 1: Validate That The Network Decoder Captures Encrypted Traffic	230
Step 2: Obtain Private Keys from Managed Servers	231
Step 3: Validate That The Private Key Cipher Suite is Supported	231
Supported Cipher Suites	231
Unsupported Cipher Suites	232
Step 4: Confirm HTTPS Parser is Enabled on Decoders	236
Step 5: Upload the Supported Private Keys to Decoders	236
Upload Multiple Premaster and Private Keys	238
Parameters for Managing Keys	241
Return Values	242
Encryption Keys	242
Premaster Key	242
Private Keys or PEM Files	242
Step 6: Reload HTTPS Parser on Decoders	243
Step 7: View Decrypted Sessions	244
Viewing Unencrypted Traffic	244
Troubleshooting	245
Decryption of Secure SMTP	245
Performance Considerations	246
TLS Certificate Hashing	248
JA3 and JA3S TLS Fingerprints	249
Perform Simultaneous Ingestion of the Encrypted and Decrypted Traffic Streams to Decoder	249
Community ID	250
Edit Decoder System Configuration	251
Enable CPU Usage Statistics for Installed Content	253
Enable Parser Mappings	254
Enable IP Address to Event Source Mapping	254
Update IP to Event Source Mapping	255
Read IP to Event Source Type Mappings	257
Edit IP to Event Source Type Mappings	257
Delete IP to Event Source Type Mappings	258
Sort the Hostname or Event Source Type	258
Import IP to Event Source Mapping Entries	259
Export IP to Event Source Mapping Entries	259
Search IP to Event Source Mapping Entries	260
Enable or Disable Lua and Flex Parsing Systems	261

Map IP Address to Service Type for Log Parsing	263
Map an IP Address to a Service Type	263
IPdevice Command	263
Map an IP Address to a Time Zone	265
tzinfo Command	265
iptmzone Command	266
Examples	266
Change the Date format	266
Missing Year Support	268
Latent log during transition to new year	269
Logs from forward time zones (or skewed clocks) just before new year transition	269
Latent logs received where a leap day cannot be successfully assigned to an appropriate leap year	269
Limitations	269
Examples	269
Obtain Log Files from a Pre-11.0 Log Decoder	271
Upload a Log File to a Log Decoder	273
Upload a Packet Capture File	274
Simultaneous Import and Capture	275
Decoder and Log Decoder References	276
Services Config View - Capture Policies Tab	277
Services Config View - Edit Capture Policies Wizard	280
Services Config View - Data Privacy Tab	283
Services Config View - Data Retention Scheduler	285
Services Config View - Feeds Tab	287
Upload Feeds Dialog	290
Services Config View - Files Tab	292
Services Config View - General Tab	294
Services Config View - Parsers Tab	303
Services Config View - Parser Mappings Tab	305
Data Export Tab	308
Services Config View - Rules Tabs	310
App Rules Tab	314
Correlation Rules Tab	318
Network Rules Tab	321
Services System View - Decoders	325

Decoder and Log Decoder Quick Setup

A basic RSA NetWitness Platform network includes at minimum Brokers, Concentrators, and Decoders. Brokers aggregate data from Concentrators, and Concentrators consume data from at least one Network Decoder or Log Decoder. The basic network may include both types of Decoders. Network Decoders are usually referred to as Decoders, and they capture network data in packet form. Log Decoders capture log data as events.

Adding a Decoder makes it visible and available for use with NetWitness Platform Administration, Live Services, and Investigate. To add a service in NetWitness Platform, you select the service type, provide service connection information, and validate that the service can be reached. The *Hosts and Services Getting Started Guide* provides the information you need to understand and install all NetWitness Platform services.

After the services are added, you need to configure each service. This is the preferred order for configuring your system:

1. Decoders
2. Log Decoders
3. Concentrators (refer to the *Broker and Concentrator Configuration Guide*)
4. Brokers (refer to the *Broker and Concentrator Configuration Guide*)

Note: A Log Decoder is a special type of Decoder, which is configured and managed in a similar way to a Decoder. Most of the information in this guide refers to both types of Decoders. "Decoder" refers to both types of Decoders. Information that applies exclusively to Network Decoders or Log Decoders is clearly identified.

Basic configuration of the Decoder involves selecting a network adapter interface and starting data capture.

In addition, you can configure each Decoder to control the type of traffic captured using rules, feeds, and parsers. Advanced configuration tasks enable additional features that are relevant to specific applications. For example, configure a 10G Decoder, create custom meta keys, or decrypt incoming packets.

The easiest way to configure all of the required Decoder and Log Decoder settings is to use the options in the NetWitness Platform user interface. For the most part, configuration is performed in the

Administration Services view ( (Admin) > Services).

Name	Licensed	Host	Type	Version	Actions
Admin Server	✓	Docker-Host	Admin Server	11.5.0.0	[Settings]
Broker	✓	Broker	Broker	11.5.0.0	[Settings]
Broker-Aggregation	✓	Broker-Aggregation	Broker	11.5.0.0	[Settings]
Broker-Foo	✓	Broker-Foo	Broker	11.5.0.0	[Settings]
Concentrator	✓	Concentrator	Concentrator	11.5.0.0	[Settings]
Concentrator-Analysis	✓	Concentrator-Anal...	Concentrator	11.5.0.0	[Settings]
Concentrator-Foo	○	Concentrator-Foo	Concentrator		[Settings]
Config Server	✓	Docker-Host	Config Server	11.5.0.0	[Settings]
Content Server	✓	Docker-Host	Content Server	11.5.0.0	[Settings]
Contexthub Server	✓	Docker-Host	Contexthub Server	11.5.0.0	[Settings]
Decoder	✓	Decoder	Decoder	11.5.0.0	[Settings]

Page 1 of 1 | Displaying 1 - 20 of 20

Administrators who feel comfortable working outside of the user interface can configure the basic parameters as well as advanced settings by editing database nodes in the Decoder node tree using the Services Explore view.

Parameter	Value
adapter.meta.enabled	on
aggregate.precache	on
assembler.parse.valve	false
assembler.pool.ratio	70
assembler.session.flush	1
assembler.session.pool	2000
assembler.size.max	32 MB
assembler.size.min	0
assembler.tcp.close	false
assembler.timeout.packet	60
assembler.timeout.session	60
assembler.voting.weights	first=1 size=1 port=1 octet=1 routable=1
capture.autostart	on
capture.buffer.size	8 MB
capture.device.params	
capture.selected	packet_mmap_eth0 (bpf)
export.cache.expire	60
export.packet.enabled	no
export.packet.local.path	/var/netwitness/decoder/export/packet

Perform Initial Quick Setup


This procedure accomplishes the initial, basic configuration of a Decoder, and starts data capture. When the basic setup is complete, the Decoder begins capturing data for the Concentrator to consume.

To configure a Decoder and start capturing data:

1. Assign a network interface for capturing data. For details, see "Select a Network Adapter" in [Configure Capture Settings](#).
2. Do one of the following:

- a. To start capture, select the Decoder and  > **View > System**. In the toolbar click

 **Start Capture**



The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes tabs for Hosts, SERVICES (selected), Event Sources, Endpoint Sources, Health & Wellness, System, and Security. Below the navigation bar, there are buttons for "Change Service", "Decoder", and "System". The main content area displays "Decoder Service Information" and "Decoder User Information".

Decoder Service Information

- Name: e92ee5cc69ef (Decoder)
- Version: 11.5.0.0 (Rev null)
- Memory Usage: 150 MB (0.47% of 32008 MB)
- CPU: 20%
- Running Since: 2020-Jun-16 04:09:33
- Uptime: 13 hours 24 minutes 9 seconds
- Current Time: 2020-Jun-16 17:33:42

Decoder User Information

- Name: admin
- Groups: Administrators
- Roles: aggregate, connections.manage, database.manage, decoder.manage, dpo.manage, index.manage, logs.manage, parsers.manage, rules.manage, sdk.content, sdk.manage, sdk.meta, sdk.packets, services.manage, storedproc.execute, storedproc.manage, sys.manage, users.manage

Session Information

Session	User	IP Address	Login Time ^	Active Queries
1451	admin	172.16.0.17-59632	2020-Jun-16 04:11:49	0
1460	admin	172.16.0.17-59634	2020-Jun-16 04:11:49	0
56627	escalateduser	172.16.0.16-49908	2020-Jun-16 14:36:12	0
56803	admin	172.16.0.16-49982	2020-Jun-16 14:38:10	0
56828	escalateduser	172.16.0.16-49982	2020-Jun-16 14:38:16	0
57517	admin	172.16.0.16-49982	2020-Jun-16 14:41:46	0
60980	admin	172.16.0.16-49982	2020-Jun-16 15:01:01	0
62026	admin	172.16.0.16-49982	2020-Jun-16 15:25:14	0

The bottom of the interface shows the RSA NetWitness Platform logo and the version number 11.5.0.0.

- b. To enable Capture Autostart, see "Configure a Decoder to Begin Capturing Data Automatically" in [Configure Capture Settings](#).


The Decoder begins capturing data for consumption by a Concentrator. For additional configuration options, refer to [Configure Common Settings on a Decoder](#) and [Decoder and Log Decoder Additional Procedures](#)

Configure Common Settings on a Decoder

This section introduces commonly used configuration settings on a Decoder with procedures and background information. After you have completed [Decoder and Log Decoder Quick Setup](#), you can refine your configuration by using parsers, feeds, and rules to limit the captured data.

Note: A Log Decoder is a special type of Decoder, which is configured and managed in a similar way to a Decoder. Most of the information in this guide refers to both types of Decoders. "Decoder" refers to both types of Decoders. Information that applies exclusively to Network Decoders or Log Decoders is clearly identified.

The following workflow illustrates commonly used settings and breaks the configuration process into four steps.




Configuration Step	Description
Configure Capture Settings	When initially setting up the Decoder, configuring the network adapter interface is required. Additional optional capture settings are available; one that is frequently used is Capture Autostart.
Enable and Disable Parsers and Log Parsers	View the parsers that have been downloaded and deployed from Live, and manage which ones are enabled or disabled.
Start and Stop Data Capture	When a Decoder starts up, it automatically begins aggregating data if Capture Autostart is enabled. When autostart is not enabled, you can start and stop data capture manually.

Configuration Step	Description
Configure Decoder Rules	Capture rules can add alerts or contextual information to sessions or logs. They can also define which data a Decoder or Log Decoder filters out. By default, no capture rules are defined when you first configure NetWitness Platform. Unless rules are specified and the rules are valid, the packets are not filtered. You can deploy the latest rules from Live as described in the <i>Live Services Management Guide</i>. You can define capture rules at any time, and you can fix rules that use invalid syntax (Fix Rules with Invalid Syntax).

Configure Capture Settings

When initially setting up the Decoder, configuring the network adapter interface is required. Additional optional capture settings are available; two that are frequently used are the Berkeley Packet Filter, and Capture Autostart.



Besides the basic network adapter interface setup, you may decide to use one of the special-purpose configurations described in [\(Optional\) Preserve VLAN Tags When Using the Packet MMAP Capture Interface](#) or [\(Optional\) Configure a Decoder to Capture Data Across All Types of Network Interfaces](#)

The rest of the capture settings have default values chosen to be effective in most cases (see a detailed list in [Services Config View - General Tab](#)). You can adjust these in some circumstances, for example, if Customer Support advises a change. You can edit the capture settings at any time.

Select a Network Adapter

The table below describes the Network Adapter settings for a Decoder. The system administrator sets the default network adapters when the Decoder is installed. Consult your System Administrator for more information.

Adapter Parameter	Description
Berkeley Packet Filter	Berkeley Packet Filters (BPF) are applied to the packet stream before the packets are copied to the Decoder adapter for analysis. This allows unwanted traffic to be efficiently discarded. However, any packets discarded are not accounted for in any Decoder statistics (capture rate, packets dropped, and packets filtered and total packets).

Adapter Parameter	Description
Capture Interface Selected	Select an adapter through which the Decoder captures packets. For the lower speed internal capture interface, use the <code>packet_mmap_,7,eth1</code> adapter, which corresponds to the monitor port located on the motherboard. There are six additional capture ports: • <code>packet_mmap_,1,lo</code> (bpf) • <code>packet_mmap_,2,eth2</code> (bpf) • <code>packet_mmap_,3,eth3</code> (bpf) • <code>packet_mmap_,4,eth4</code> (bpf) • <code>packet_mmap_,5,eth5</code> (bpf) • <code>packet_mmap_,8,ALL</code> (bpf) There are three wireless capture services available: • <code>packet_netmon_</code> (Microsoft Netmon) • <code>packet_mac80211_</code> (Linux mac80211) • <code>packet_airport_</code> (Mac OS X AirPort)
Capture Interface Selected for Log Decoder	The following capture service is available: • <code>log_events,Log Events</code>

To configure the network adapter on a Decoder:

1. Go to  (Admin) > **Services**.
2. In the **Administration Services view**, select the Decoder and  > **View** > **Config**.
The Services Config view is displayed with the General tab open.

Decoder Configuration	
Name	Config Value
Adapter	
Berkeley Packet Filter	
Capture Interface Selected	
Cache	
Cache Directory	/var/netwitness/decoder/cache
Cache Size	4 GB
Capture Settings	
Assembler Maximum Size	32 MB
Assembler Minimum Size	0
Assembler Session Flush	1
Assembler Session Pool	250000
Assembler Timeout Packets	60
Assembler Timeout Session	60
Capture Autostart	☐
Capture Buffer Size	64 MB

3. In the **Capture Interface Selected** field, select the network adapter that best suits the Decoder.
4. To save the changes, click **Apply**.
5. If necessary to put the changes into effect, navigate back up to the **Administration Services** view, select the Decoder, and select   > **Restart**.

Configure a Decoder to Begin Capturing Data Automatically

1. Go to  (Admin) > **Services**.
2. In the **Administration Services** view, select the Decoder and   > **View** > **Config**.
The Services Config view is displayed with the General tab open.

Decoder Configuration	
Name	Config Value
Adapter	
Berkeley Packet Filter	
Capture Interface Selected	
Cache	
Cache Directory	/var/netwitness/decoder/cache
Cache Size	4 GB
Capture Settings	
Assembler Maximum Size	32 MB
Assembler Minimum Size	0
Assembler Session Flush	1
Assembler Session Pool	250000
Assembler Timeout Packets	60
Assembler Timeout Session	60
Capture Autostart	☐
Capture Buffer Size	64 MB

3. Under **Capture Settings**, select the **Capture Autostart** checkbox.
4. To save the changes, click **Apply**.
5. If necessary to put the changes into effect, navigate back up to the **Administration Services view**, select the Decoder, and select   > **Restart**.

Configure Optional Capture Settings

1. Go to  (Admin) > **Services**.
2. In the **Administration Services view**, select the Decoder and   > **View > Config**.
The Services Config view is displayed with the General tab open.

Decoder Configuration	
Name	Config Value
Assembler Session Flush	1
Assembler Session Pool	250000
Assembler Timeout Packets	60
Assembler Timeout Session	60
Capture Autostart	☐
Capture Buffer Size	64 MB
Parse Maximum Bytes	128 KB
Parse Minimum Bytes	1 KB
Parse Threads	0
Database Max File Sizes	
Meta File Size	auto
Packet File Size	auto
Session File Size	auto
Hash	
Hash Directory	

3. If you want to apply a system-level filter to the packet stream before the packets are copied to the Decoder adapter for analysis, configure the Berkeley Packet Filter as described in [\(Optional\) Configure System-Level \(BPF\) Packet Filtering](#).
4. In the **Capture Settings** sections, review the default values. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change. See [Services Config View - General Tab](#) for an explanation of these settings.
5. In the **Database Max File Sizes** section, review the default values. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change. See [Services Config View - General Tab](#) for an explanation of these settings.
6. In the **Hash** section, define a directory for hash files if you are using this feature. See [Services Config View - General Tab](#) for an explanation of these settings.

(Optional) Configure System-Level (BPF) Packet Filtering

You can use Berkeley Packet Filters to control which packets and logs are processed by a Decoder.

Berkeley Packet Filters (BPF) are applied to the packet stream before the packets are copied to the Decoder adapter for analysis. This allows unwanted traffic to be efficiently discarded. These discarded packets are not accounted for in any Decoder statistics (capture rate, packets dropped, and packets filtered and total packets).

The Decoder also supports system-level packet filtering defined using `tcpdump/libpcap` syntax. Specifying a `Libpcap` filter can efficiently reduce packet volume based on Layer 2 - Layer 4 attributes. A `Libpcap` filter is appropriate for use when a Decoder is receiving a traffic volume that is placing a load against the physical resources of the platform. In this scenario, the Decoder may consistently drop packets and have a large number of capture pages available (`/decoder/stats/capture.pagefree` is high).

The following is an example of a `libpcap` filter to keep only packets that do not have both source and destination addresses in the 10.21.0.0/16 subnet.

```
not (src net 10.21.0.0/16 and dst net 10.21.0.0/16)
```


For a full reference of the `Libpcap` filter syntax, see the main pages for:

- `tcpdump` (http://www.tcpdump.org/tcpdump_man.html).
- `pcap-filter` (<http://www.unix.com/man-page/FreeBSD/7/pcap-filter/>).

To add a system-level Berkeley Packet Filter:

1. Go to  (Admin) > **Services**.

- In the Administration Services view, select a Decoder service and  > **View > Config**. The Services Config view is displayed with the General tab open.



The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes tabs for Hosts, Services, Event Sources, Endpoint Sources, Health & Wellness, System, and Security. The 'Services' tab is active, and the 'Config' sub-tab is selected. The 'General' configuration tab is open, displaying three main sections:

- System Configuration:** A table with columns 'Name' and 'Config Value'.

Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☒
SSL Port	56004
Stat Update Interval	1000
Threads	20
- Decoder Configuration:** A table with columns 'Name' and 'Config Value'.

Name	Config Value
Adapter	
Berkeley Packet Filter	
Capture Interface Selected	packet_mmap_eth0 (bpf)
Cache	
Cache Directory	/var/netwitness/decoder/cache
Cache Size	4 GB
Capture Settings	
Assembler Maximum Size	32 MB
Assembler Minimum Size	0
Assembler Session Flush	1
Assembler Session Pool	2000
Assembler Timeout Packets	60
Assembler Timeout Session	60
- Parsers Configuration:** A table with columns 'Name' and 'Config Value'.

Name	Config Value
ALERTS	Enabled
DHCP	Enabled
DNS	Enabled
Entropy	Disabled
FeedParser	Enabled
FTP	Enabled
GeoIP2	(Mixed)
GTalk	Enabled
H323	Enabled
HTTP	Enabled
HTTP2	Enabled
HTTPS	Enabled
IRC	Enabled
JSNormalize	Enabled
MAIL	Enabled
NETBIOS	Enabled
NETWORK	Enabled
NFS	Enabled
NNTP	Enabled
OPENAPPID	Enabled
PGP	Enabled

An 'Apply' button is located at the bottom right of the configuration area. The footer shows 'RSA NETWITNESS PLATFORM' and version '11.5.0.0'.

- In the **Decoder Configuration Section**, under **Adapter**, click in the field next to **Berkeley Packet Filter**.
- Type only one filter in the field. If you want to filter multiple items, join multiple expressions using `and`. Several examples are provided below.
The user interface validates input at the time you enter your filter string.
- To save the filter, click **Apply**.
If the syntax is correct, a confirmation message is displayed.
If the syntax is incorrect, a **Packet filter is not valid** message is displayed and a corresponding log message will follow in the log messages on the Decoder:

```
164474800      2020-May-01 19:03:08      warning      Decoder      Failed to
parse filter 'example_badrule': syntax error
```
- To activate the filter, you must stop and start capture on the Decoder:
 - Change the **Config** view to the **System** view.
 - Click **Stop Capture**.
 - Click **Start Capture**.
The active filter will be displayed in the Decoder logs.

Examples

These are several filter examples:

- Drop packets to or from any address in the 10.21.0.0/16 subnet:
`not (net 10.21.0.0/16)`
- Drop packets that have both source and destination addresses in the 10.21.0.0/16 subnet:
`not (src net 10.21.0.0/16 and dst net 10.21.0.0/16)`
- Drop packets that are from 10.21.1.2 or are headed to 10.21.1.3.
`not (src host 10.21.1.2 or dst host 10.21.1.3)`
- Combine both IP and HOST:
`not (host 192.168.1.10) and not (host api.wxbug.net)`
- Drop all port 53 traffic, both TCP & UDP:
`not (port 53)`
- Drop only UDP port 53 traffic:
`not (udp port 53)`
- Drop all IP protocol 50 (IPSEC) traffic:
`not (ip proto 50)`
- Drop all traffic on TCP ports 133 through 135.
`not (tcp portrange 133-135)`

The following filters combine some of the above to demonstrate how to put multiple directives into one filter:

- Drop any port 53(DNS) traffic sourced from 10.21.1.2 or destined to 10.21.1.3.
`not (port 53) and not (src host 10.21.1.2 or dst host 10.21.1.3)`
- Drop any traffic using IP proto 50 or port 53 or any traffic from net 10.21.0.0/16 destined to net 10.21.0.0/16
`not (ip proto 50 or port 53) or not (src net 10.21.0.0/16 and dst net 10.21.0.0/16)`

Caution: The use of parentheses can have a large and potentially disruptive effect on the use of Packet Filters. As a best practice, keep "not" operations outside of parentheses and always test your rules before deploying them. Failure to properly format your rules (despite input validation) can cause a packet filter to drop ALL traffic or behave in other unexpected ways. This is due to the way packet Libpcap filters work and is not the result of any logic within NetWitness Platform software.

Testing

BPF filters can and should be tested using either `tcpdump` or `windump` to ensure that they will provide the expected behavior before implementing them. This example shows a test of a filter using `windump`:

```
windump -nni 2 not (port 53 or port 443) or not (ip proto 50)
```

Conversions

If for the sake of performance, you have decided that an existing network rule filter would be better running as a System-Level Packet Filter, you can convert it. There are a few things to remember when doing conversions.

- `&&` becomes `and`
- `alias.ip` becomes `host` if a single host or `net` if a network.
- `ip.src` becomes `src host` if a single host or `src net` if a network.
- `ip.dst` becomes `dst host` if a single host or `dst net` if a network.
- Use CIDR notation when listing a network (that is, `10.10.10.0/24`).
- `||` becomes `or`
- `!` becomes `not`
- Multiple rules must be joined with `and`.

The manual for TCPDump also gives examples of filters and strings that can be used:

http://www.tcpdump.org/tcpdump_man.html

Additionally, the following site provides an excellent reference for BPF-style packet filters:

<http://biot.com/capstats/bpf.html>

Caution: If you are capturing `vlan` tagged packets, above standard `bpf` filter may not work. For example, if you use `not (udp port 123)` to filter `vlan` tagged NTP traffic on `udp` port 123, it will not work. This is because the `bpf` filter machinery is simple and does not account for protocols not referenced in the rule. So the OS executing the `bpf` filter will look for the `udp port` values at the byte offset they would occur in a standard Ethernet/`udp` packet; but the optional `vlan` tag fields in the Ethernet header pushes those values by 4 bytes, thus the `bpf` filter rule will fail. To fix it, you need to change the `bpf` filter to: `not (vlan and udp port 123)`.

(Optional) Configure a Decoder to Capture Data Across All Types of Network Interfaces

The `packet_mmap_`, `ALL` adapter is capable of capturing across all types of network interfaces at the same time. For example, this can include things like physical network interfaces over different media types and tunnel interfaces.

The default behavior of the `ALL` adapter is to capture from all interfaces from the system, except for the hard-coded defaults of `lo`, `eth0`, and `eml`.

You can select any subset of the capture interfaces by editing the Decoder configuration node `/decoder/config/capture.device.params` to include an `interfaces=` parameter. The `interfaces` parameter contains a comma-separated list of interfaces that are used for capture. Instead of using all interfaces for capture, only the specified interfaces are used.


For example, if you want to force capture on interfaces `em1`, `em2`, and `em4`, and ignore `em3`, you can select the `packet_mmap_`, `ALL` adapter, and then add this line to `capture.device.params`:
`interfaces=em1,em2,em4`

Note: Decoder automatically performs decapsulation of Virtual Extensible LAN (VXLAN) protocol from network traffic on UDP-4789 so that parsing can take place on the decapsulated Ethernet frames. During the decapsulation, the VXLAN ID is extracted and stored in the VXLAN meta key. By default, the VXLAN ID extraction is enabled. To disable, navigate to the `vlanGre` parser in the Parser Configuration settings. Click the drop-down list for `vlan` in the **Config Value** column and select **Disabled**. Each parser is configurable in the [Services Config View - General Tab](#). The Parser Configuration panel provides a way to enable or disable parsers to use on Decoders in addition to limiting the metadata that the parser creates.

Note: Using the `interfaces` parameter to select `eth0`, `lo`, or `em1` overrides the default behavior, which is to drop traffic from those ports.

To configure the `packet_mmap_`, `ALL` adapter to capture from specific interfaces instead of all interfaces:


1. Go to  (Admin) > **Services**, select the Decoder service and  > **View** > **Config**.
2. In the **Services Config view**, set **Capture Interface Selected** to `packet_mmap_`, `ALL` adapter.



The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes tabs for Hosts, Services, Event Sources, Endpoint Sources, Health & Wellness, System, and Security. The 'Services' tab is selected, and the 'Decoder' service is chosen. The 'Config' view is open, showing the 'General' tab. The 'System Configuration' section includes fields for Compression, Port, SSL FIPS Mode, SSL Port, Stat Update Interval, and Threads. The 'Decoder Configuration' section includes fields for Adapter, Berkeley Packet Filter, Capture Interface Selected (with a dropdown menu open showing options like 'packet_mmap_eth0 (bpf)', 'packet_mmap_lo (bpf)', and 'packet_mmap_ALL (bpf)'), Cache, and Capture Settings. The 'Parsers Configuration' section on the right lists various parsers with their status (Enabled or Disabled). At the bottom, there is an 'Apply' button.

3. Click **Apply**, and then restart the Decoder service.

- Go to  (Admin) > **Services**, select the Decoder service, and click  > **Config** > **Explore**.
- In the Services Explore view, select **decoder** > **config**.



The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes tabs for Investigate, Respond, Users, Hosts, Files, Dashboard, and Reports. Below this is a secondary navigation bar with categories like HOSTS, SERVICES, EVENT SOURCES, ENDPOINT SOURCES, HEALTH & WELLNESS, SYSTEM, and SECURITY. The main content area is titled '115Decoder - Decoder' and shows a list of configuration items under the 'config' tab. The left sidebar shows a tree view of the service structure, including connections, database, decoder, config, devices, parsers, stats, index, logs, rest, sdk, services, storedproc, sys, and users. The main table lists various configuration parameters and their current values.

Path	Value
/decoder/config	115Decoder - Decoder
Aggregate Precache (aggregate.precache)	on
Parse Blowoff Valve (assembler.parse.valve)	false
Assembler Pool Ratio (assembler.pool.ratio)	70
Assembler Session Flush (assembler.session.flush)	1
Assembler Session Pool (assembler.session.pool)	1000000
Assembler Maximum Size (assembler.size.max)	32 MB
Assembler Minimum Size (assembler.size.min)	0
Parse on TCP Close (assembler.tcp.close)	false
Assembler Timeout Packets (assembler.timeout.packet)	60
Assembler Timeout Session (assembler.timeout.session)	60
Assembler Voting Weights (assembler.voting.weights)	first=1 size=1 port=1 octet=1 routable=1
Capture Autostart (capture.autostart)	on
Capture Buffer Size (capture.buffer.size)	128 MB
Capture Device Parameters (capture.device.params)	
Capture Interface Selected (capture.selected)	PFRINGZC,em3 (bpf);PFRINGZC,em4 (bpf)
Write Capture Port Meta (captureport.meta.enabled)	on
Seconds Before Cache Expiring (export.cache.expire)	60
Packet Remote Export (export.packet.enabled)	no
Packet Local Export Path (export.packet.local.path)	/var/netwitness/decoder/export/packet
Maximum Packets Per File (export.packet.max)	0
Packet Remote Export Path (export.packet.remote.path)	/var/netwitness/decoder/export/packet

- Click in the values column next to `capture.device.params`, type **interfaces=em1,em2,em4**, and press **Enter**.

Path	Value
/decoder/config	Decoder
adapter.meta.enabled	on
aggregate.precache	on
assembler.parse.valve	false
assembler.pool.ratio	70
assembler.session.flush	1
assembler.session.pool	2000
assembler.size.max	32 MB
assembler.size.min	0
assembler.tcp.close	false
assembler.timeout.packet	60
assembler.timeout.session	60
assembler.voting.weights	first=1 size=1 port=1 octet=1 routable=1
capture.autostart	on
capture.buffer.size	8 MB
capture.device.params	interfaces=em1,em2,em4
capture.selected	packet_mmap_eth0 (bpf)
export.cache.expire	60
export.packet.enabled	no
export.packet.local.path	/var/netwitness/decoder/export/packet
export.packet.max	0
export.packet.remote.path	/var/netwitness/decoder/export/packet
export.packet.size.max	1 GB
export.session.enabled	no
export.session.format	avro
export.session.local.path	/var/netwitness/decoder/export/session
export.session.max	0
export.session.meta.fields	*
export.session.remote.path	nfs://magr/saw/sessions

The change goes into effect immediately; only traffic on em1, em2, and em4 interfaces is captured.

(Optional) Configure Meta-Only Decoders

You can configure Decoders so that packets and logs can be processed, and then dropped before they are written to disk. This is called a Meta-Only Decoder, which uses the Meta-Only license, and can save storage space (however, analysts cannot reconstruct events in Investigate if you use this option). The configuration option `/decoder/config/packet.write.disabled` controls this feature. If this option is set to `true`, all packets are dropped after parsing, so they are never written to the database. This applies to both Log and Network Decoders. The ingested logs and packets flow through the system normally so that parsing and other operations are not impacted. The default setting is `false`, which preserves the normal behavior of writing packets to disk.

Note: You must purchase the Meta-Only license before you can realize the full benefit of using this option. For information about purchasing licenses, see the *Licensing Management Guide for RSA NetWitness Platform*. Go to the [Master Table of Contents](#) to find all RSA NetWitness Platform 11.x documents.

To configure a Meta-Only Decoder where packets and logs are parsed and not written to disk:

1. Go to  (Admin) > **Services** and select a Decoder.
2. Click  > **View** > **Explore**, and in the left panel, expand **decoder** and click **config**.

3. In the right pane, go to `packet.write.disabled`, and change the value from `false` to `true`.

Note: You can make this configuration update while capture is running, but the update does not take affect until capture is restarted.

For an example of how to use the Meta-Only license to apply centrally-managed capture policies across your Network Decoders, see [\(Optional\) Configure Selective Network Data Collection](#)

(Optional) Configure Selective Network Data Collection

Selective network data collection gives administrators the ability to apply centrally managed capture policies across their Network Decoders. This results in better use of service resources, including hard drive space, which leads to more predictable costs and lessens the burden of managing multiple services. Administrators can determine which traffic is stored and how it is stored by using policies. Each policy contains a list of supported base protocols and definitions for handling any other protocols that are detected.

Note: Administrators can use this feature to create policies for metadata-only while using the cost-effective Meta-Only license, without having to pay for additional storage or throughput licensing. To use this license, work with your RSA Customer Support Representative to purchase the license, and then enable the functionality as described in [\(Optional\) Configure Meta-Only Decoders](#).

The administrator can choose to deploy predefined policies that capture:

- All base and other protocols (Full Capture - All Protocols)
- Only metadata on all base and other protocols (Capture Meta Only - All Protocols)
- Only metadata on all base protocols and drop all other protocols (Capture Meta on Base Protocols, Drop all other protocols)
- All base protocols and only metadata on all other protocols (Full Capture on Base Protocols, Meta only on all other protocols)

The predefined policies are not configurable. The only way you can edit predefined policies is by assigning services (such as Decoders) to deploy them.

Administrators can create custom policies to give further control over the deployment. A base set of protocols are available for alterations by the administrator, allowing you to choose what level of capture you prefer on a per-protocol basis. If you are only making slight changes, a good start for customization is to clone one of the predefined policies and alter it. These centrally-managed policies are then applied to services (Network Decoders) to allow handling multiple use cases across your environment.

Here are some example use cases and how an administrator would deploy these.

Use Case 1: Predefined Policy for Internal Assets

To gain further visibility into internal lateral traffic going East-West through the environment, a Network Decoder is deployed in the interior of the network to capture this traffic. Applying full packet capture for this segment may generate too much data to manage at a reasonable cost point. Instead, the predefined "Capture Meta Only – All Protocols" policy can be applied to achieve the same parsing and detection while only saving the metadata, limiting the amount of storage required of the solution. In this situation, the only steps you must complete is to choose the predefined policy, edit it, and then enable it by assigning it to the appropriate internal services (Decoders). For an example of how to set up this policy, see [Use Predefined Policies](#).

Use Case 2: Custom Policy for DMZ Assets

To protect the public-facing web servers, a Network Decoder is deployed into the demilitarized zone (DMZ). The network communications between these web servers and other network devices are a subset of the base protocols. To make sure you can harvest forensic evidence for any communications related to these high-value assets, you would want to do full packet capture on HTTP, SMTP, FTP, and DNS protocols. You still want visibility into any other traffic these servers are engaged in, so you would prefer metadata generated for all other protocols. The suggested steps to create this custom policy would have you choose the predefined policy "Full Capture on Base Protocols, Meta only on all other protocols" and then clone it to adjust it to the protocols aligned with the DMZ scenario. During the clone process, the name and description can be updated to identify the new purpose of the policy. The define policy step allows all the base protocol rules inside the policy except HTTP, SMTP, FTP, and DNS to be altered from "Collect All" to "Collect Meta Only" to achieve the desired policy. The last steps are assigning the policy to the appropriate DMZ services (Decoders). For an example of how to set up this policy, see [Create New Policies from Predefined Ones](#).

Note: When a Decoder is being managed by a selective collection policy, it is possible that the Decoder specifications in the policy could be overwritten if the Decoder configuration is edited outside of the policy. To try and avoid this situation, the following warning message is displayed at the top of the Parsers section of the Decoder Configuration page: This Decoder is being managed by the following capture policy: <policy-name>. Some parsers are required by the capture policy, and changes made here may overwrite that policy. Review of the capture policy is recommended.

Another warning message is displayed at the top of the Application Rule section: This Decoder is being managed by the following capture policy: <policy-name>. Some application rules are required by the capture policy, and changes made here may overwrite that policy. Review of the capture policy is recommended.

The following sections describe configuring and using selective network data collection.

- [Use Predefined Policies](#)
- [Create New Policies from Predefined Ones](#)
- [Create Custom Policies](#)
- [Unpublish Policies](#)
- [Delete Policies](#)
- [Supported Protocols for Selective Network Data Collection](#)

For detailed information about the Selective Collection user interface, see:

- [Services Config View - Capture Policies Tab](#)
- [Services Config View - Edit Capture Policies Wizard](#)

Use Predefined Policies

This procedure describes how to use the predefined policies. You can use this procedure to configure a policy for internal assets. For example, as described in [Use Case 1: Predefined Policy for Internal Assets](#), to gain further visibility into internal lateral traffic going East-West through the environment, you could deploy a Network Decoder in the interior of your network to capture the traffic. Applying full packet capture for this segment may generate too much data to manage at a reasonable cost point. Instead, use the predefined "Capture Meta Only – All Protocols" policy to achieve the same parsing and detection while only saving the metadata, limiting the amount of storage required for this solution.

In this situation, the only steps you must complete is to select the "Capture Meta Only – All Protocols" predefined policy to edit it, enable it by assigning it to the appropriate Decoders, and then publish the policy to start collecting the data.

To work with a predefined policy:

1. In the NetWitness Platform user interface, go to  (Configure) and select the **CAPTURE POLICIES** tab.

Notice that the PUBLICATION STATUS column lists each of the predefined policies with the status of **Unpublished**. The only way to edit a predefined policy is to assign Decoders to apply the policy. If changes are made to the Decoder assignments but are not published, the status changes to **Unpublished Edits**, the date the change was made is displayed in the POLICY UPDATED column, and the user name of the person who made the change is displayed in the UPDATED BY column. In predefined policies, the CREATED BY field is always **system**.

The following figure shows the statuses of the policies.

2. Select **Capture Meta Only - All Protocols**, and then click **Edit**.
The editing view is displayed with the **Identify Policy** option active.
3. Because this is a predefined policy, you cannot change the policy name or description. Click **Next**.
The Define Policy: Capture Meta Only - All Protocols page is displayed.
4. Because this is a predefined policy, you cannot edit this page. Review the rules on this page to understand the actions that will be enabled when you publish this policy. Click **Next**.
The Assign Policy page is displayed.
5. Select the Decoders that you want to collect from, and then click **Save and Publish**.
The publication status moves to **Updating**. It changes to **Published** when the policy has been successfully deployed to all the Decoders that you selected. However, the Decoders that are available start gathering data immediately, before the policy is displayed as **Published**, so even if not all of the Decoders are available (and the publication status is not **Published**), data will begin to be collected on the Decoders that are available. When the status moves to **Published**, all the services are displayed in the Service Assignment column. You can hover over the names of the Decoders to see their full names, as shown in the following figure.

Create New Policies from Predefined Ones

This procedure describes the steps for creating a new policy from a predefined one, and provides an example of how you might use this policy.

For example, as described in [Use Case 2: Custom Policy for DMZ Assets](#), suppose a Network Decoder has been deployed into the demilitarized zone (DMZ) to protect public-facing web servers. The network communications between these web servers and other network devices are a subset of the base protocols. To make sure they can harvest forensic evidence for any communications related to these high value assets, you would want to get full packet capture on HTTP, SMTP, FTP, and DNS protocols. To provide visibility into any other traffic these servers are engaged in, you would want to generate metadata for all other protocols.

To configure a policy that would achieve the results described here, you would select the predefined policy "Full Capture on Base Protocols, Meta only on all other protocols," clone it, and then edit it to align the protocols with the DMZ scenario. During the clone process, you can update the name and description to identify the new purpose of the policy. In the step where you define the policy, you can change all the base protocol rules inside the policy (except HTTP, SMTP, FTP, and DNS) from "Collect All" to "Collect Meta Only" to achieve the desired results. The last steps assign the policy to the appropriate DMZ Decoders.

1. In the NetWitness Platform user interface, go to  (Configure) and select the **CAPTURE POLICIES** tab.
2. Select **Full Capture on Base Protocols, Meta only on all other protocols** and then click **Clone**. The Collection Policies view is displayed with the cloned policy at the bottom of the list, identified by appending -1 to the policy name.
3. Select the cloned policy and click **Edit**.
The editing view is displayed with the **Identify Policy** option active.
In this example, we are naming this policy "DMZ Assets" and providing the description "This policy is defined to protect public-facing web servers."
4. Click **Next**. The Define Policy view is displayed, with a list of rules you can use to define the policy. You can select the following options for rule actions:
 - **Collect Meta Only**: Collect metadata
 - **Drop All**: Drop metadata and network packets
 - **Collect All**: Collect metadata and network packetsTo collect the targeted data for the DMZ Assets policy, we are going to change all the base protocol rules inside the policy (except HTTP, SMTP, FTP, and DNS) from **Collect All** to **Collect Meta Only**, as shown below.

5. Click **Next**.

The Assign Policy page is displayed.

6. A Decoder can only be assigned to one policy. Decoders that are not available are greyed-out, because they are already assigned to other policies. Select the Decoders to collect the targeted data, and then click **Save and Publish**. The publication status moves to **Updating**. It changes to **Published** when the policy has been successfully deployed to all the Decoders that you selected. However, the Decoders that are available start gathering data immediately, before the policy is displayed as **Published**, so even if not all of the Decoders are available (and the publication status is not **Published**), data will begin to be collected on the Decoders that are available. Your new policy is displayed in the Collection Policies view, with the services displayed in the Service Assignment column.

Create Custom Policies

You can create a custom Selective Network Decoder policy, where you can select specific protocols on which to collect data, as well as the rules that apply to them.

1. In the NetWitness Platform user interface, go to  (Configure) and select the **CAPTURE POLICIES** tab.
2. Click **Create New**.
The editing view is displayed with the **Identify Policy** option active.
3. Enter a name and description for your policy, and then click **Next**.
The Define Policy view is displayed.
4. Define the actions for each of the protocols. The options for **RULE ACTION** are:
 - **Collect Meta Only**: Collect metadata
 - **Drop All**: Drop metadata and network packets
 - **Collect All**: Collect metadata and network packets
5. Click **Next**.
The Assign Policy view is displayed.

6. A Decoder can only be assigned to one policy. Decoders that are not available are greyed-out, because they are already assigned to other policies. Select the Decoders to collect the targeted data, and then click **Save and Publish**. The publication status moves to **Updating**. It changes to **Published** when the policy has been successfully deployed to all the Decoders that you selected. However, the Decoders that are available start gathering data immediately, before the policy is displayed as **Published**, so even if not all of the Decoders are available (and the publication status is not **Published**), data will begin to be collected on the Decoders that are available.

Your new policy is displayed in the Collection Policies view, with the services displayed in the Service Assignment column.

Verify Published Policies

When a policy has been published, the policy is applied to the Decoders that were designated in the policy assignment step. The Decoders that are assigned to the policy start collecting the data defined in the policy. There are a few ways to validate the policy is functioning as expected.

View Publication Status

Publication status is indicated in the Capture Policies view in the PUBLICATION STATUS column. You can also see which Decoders are assigned to the policy in the SERVICE ASSIGNMENT COLUMN.

The publication statuses are:

- **Unpublished** - the rule actions of the policy have been defined and Decoders may have been assigned to the policy, but the policy has not been deployed to any Decoders.
- **Unpublished Edits** - policies that have been updated but are not deployed to Decoders. Clicking **Save and Close** after making edits keeps them from automatically being published, allowing the administrator to wait until a specific outage window to publish them. This would display as **Unpublished** (for policies not previously published to Decoders) or **Unpublished Edits** (for policies that have been previously published to Decoders, but the updates have not been deployed to Decoders).
- **Published** - the policy has been successfully deployed to all the Decoders that you selected. However, the Decoders that are available start gathering data immediately, before the policy is displayed as Published, so even if not all of the Decoders are available (and the publication status is not Published), data will begin to be collected on the Decoders that are available.
- **Failed** - policies that failed to execute, for example, if a Decoder is offline or a system is down.

You can also select a published policy, and then click **Edit** to verify the Decoders that are assigned to the policy.

Verify Policies in Admin > Services

1. As an administrator, log in to the NetWitness Platform user interface to view the Decoders that have been assigned to the policy.

2. Go to  (Admin) > **Services** > <decoder name> >   (Actions) > **View** > **Config**. On the General tab, in the right panel, there is a list of parsers that are installed on the Decoder by default. If a policy is administering this Decoder and requires any of those parsers to function, a message is displayed as shown below.
3. On the **App Rules** tab, a similar message indicates that some of the rules are required for the policy to function. You can search for "selective collection" to find the specific rules required, as the rule names contain that information.

Note: In Investigate, you must have the **Alert** session option selected for the selective collection policy to be displayed.

By default, the selective collection rules are added to the bottom of the rule set when the policy is applied. However, an administrator can directly affect the policy by altering the selective policy rules or by moving conflicting rules (rules with the same services) above the selective collection rules. In both cases, those changes take effect when they are saved locally on the individual Decoders, until the next time the policy is published from the central configuration page, when that policy will override changes made individually on the Decoders.

Note: There currently is no feedback in the centralized configuration page denoting that the configuration related to the policy has been altered locally on the Decoders. To avoid this situation, manage all the policy rules from the centralized configuration page. If that is not possible, as additional application rules outside of collection policies may be required, follow standard security practices, including separation of duties, to limit people who need access to alter application rules and general parsers.

Verify Policies in the Investigate View

You can further verify that the policy is working as expected (the protocols are being collected as published) in the Investigate view. In the centralized configuration page for selective collection, the administrator can get a translation of what the query parameter would be for each individual protocol when viewing the rules inside the policy.

For example, in  (Configure) > **Capture Policies** > <select a policy> > **Edit** > **Define Policy**, if you hover over the DNS protocol name, `service=53` is displayed.

To determine if the appropriate action has been taken for DNS traffic, in **Investigate** > **Events** you can add the query parameter (for example, `service=53`) as a filter in the query bar and run the query against the Broker or Concentrator the Network Decoder is associated with for the time the policy has been active on the Network Decoder. To make sure you are only seeing the traffic from the Decoders that you want, you can add the Decoder Source, or `did` meta key as another filter in the query bar.

After running the query, if there are any related results, those events are displayed in the table (unless you applied an action to drop them).

If you chose the action to `collect all` you can click on one of the events and review the metadata and the payload information for that packet session.

If you chose the action to `collect meta only`, you can click on one of the events and review the metadata, however, a message would be displayed instead of the packet payloads, as shown in the following figure.

To expedite this process, you can also take the application rules from one of the Network Decoders the policy is applied to and use it as the filter in the **Investigate > Events** query bar.

The following figure shows an entire `collect meta only` configuration by using the application rule shown here, created on the Decoder (in [Verify Policies in Admin > Services](#)):

```
(service=0,20,21,22,23,25,53,67,69,80,110,119,137,139,161,443,520,1433,1521,1719,1720,2000,2049,5060,5222,6667) truncate:
```

Note: If the findings do not align with what is expected from the policy, make sure that there are not other rules on the individual Decoders that might be conflicting with the selective collection policy rules. In addition, make sure when checking in Investigate that you are looking at only the traffic associated with the Decoders related to the policy.

Troubleshooting Policy Deployments That Fail

Policy deployment can fail for the following reasons:

- A Decoder service assigned to the policy is not running
- A Decoder host system is down when you try to deploy the policy

When a policy deployment fails, the Failed icon is displayed in the PUBLICATION STATUS and SERVICE ASSIGNMENT columns as shown in the following figure.

Use the following options to research publication failures:

Check the status of the Decoders

Check to see if the Decoder system is down, or if the Decoder service has stopped capturing data.

1. Select the policy, click **Edit**, and go to the Assign Policy page to determine which Decoders are assigned to the policy,
2. Go to  (Admin) > **Services**, select the Decoder, click  > **View > System** to investigate the status of the Decoder.

Review the content-server.log file

You can review the content in this file to troubleshoot errors. This file is located in the following directory on the NW Server host:

```
/var/log/netwitness/content-server/content-server.log
```

The following example shows a content-server log snippet that is displayed when publishing a policy fails:

```
2020-07-08 16:22:45,287 [ nioEventLoopGroup-2-2]
INFO      DataAccess|Finished publishing policy 'Full Capture on Base
Protocols, Meta only on all other protocols-1' outcome:FAILED
```

This is an example of a content server log snippet for successfully publishing a policy:

```
2020-06-23 17:40:50,373 [ nioEventLoopGroup-2-7]
INFO      DataAccess|Finished publishing policy 'Policy 1' outcome:SUCCESS
```

Review /var/log/messages on Decoder Systems

You can view publication status in /var/log/messages on a Decoder system.

The following is an example of successful publication in /var/log/messages:

```
Jun 15 18:35:17 rsa-nw-1141-Decoder-1 NwDecoder[8260]: [Rules] [audit] User
admin (session 26182, 10.237.169.43:55380) has added application rule
'selective-collection:meta-only' at position 38
```

Unpublish Policies

You can unpublish a policy by unassigning the Network Decoders from the policy, which preserves the policy but makes it inactive so that the data it defines is not being collected. You can then change or remove a portion of policy, or remove the entire policy from the Decoders to which they are assigned.

To unpublish a policy:

1. In the NetWitness Platform user interface, go to  (Configure) and select the **CAPTURE POLICIES** tab.
2. Select the published policy and click **Edit**.
 - a. To change the name or description of the policy, update the policy name or description in the Identify Policy view.
 - b. To change the protocols that are collected, click **Next** and then edit the rule actions in the Define Policy view.
 - c. To unassign or reassign Decoders to this policy, click **Next**.
 - To fully unpublish this policy, deselect all the Decoders and then click **Save and Publish**.
 - To change the Decoders that are assigned to this policy, select or deselect the Decoders as needed and then click **Save and Publish**.

When the changes to the policy status have been completed and the Decoders updated, the policy status is displayed in green text as **Published** on the Collection Policies tab. If you unassigned all the Decoders without assigning any new ones, the publication status is **Unpublished Edits**.

Delete Policies

You can only delete customized policies. Predefined policies can be unpublished, but cannot be deleted. You can only delete one policy at a time. When you delete a policy, any Decoders that were assigned to the policy immediately stop applying the rules in the policy.

For information about how to unpublish a policy, see [Unpublish Policies](#).

To delete a policy:

1. In the NetWitness Platform user interface, go to  (Configure) and select the **CAPTURE POLICIES** tab.
2. Select a customized policy and click **Delete**.
The policy is removed from the system.

Caution: If you delete a policy that is published to Decoders, a warning is displayed. If you click **Accept** in the warning, that policy is removed from the Decoders, the collection status is changed, and the policy is deleted.

Supported Protocols for Selective Network Data Collection

The protocols in this list are supported for selective network data collection. These protocols have descriptive tool tips in the user interface. You can choose how to handle each of these protocols individually by dropping, capturing metadata only, or capturing all data related to them. There is also an additional rule for each policy that handles all other unidentified protocols not included in this list. The policies apply the collect or drop actions based on protocol characteristics, regardless of the port that they are captured on.

Category	Protocol	Description
Database	TNS	Transparent Network Substrate Oracle database networking protocol messages
	TDS	Tabular Data Stream for SQL server communications
Email	SMTP	Simple Mail Transport Protocol messages
	POP3	Post Office Protocol v3 messages (Email clients use to retrieve mail from server)
File	SMB	Server Message Block (or CIFS) network communication protocol file sharing communications
	NFS	Network File System protocol messages
	FTP	File Transfer Protocol messages
	TFTP	Trivial File Transfer Protocol messages
Networking	SNMP	Simple Network Management Protocol messages
	DNS	Domain Name System messages
	DHCP	Dynamic Host Configuration Protocol messages
	NETBIOS	Network Basic Input/Output System networking protocol messages

Category	Protocol	Description
Remote Access	SSH	Secure Shell secure remote connections
	TELNET	Teletype network remote interface protocol messages
	SCCP	Cisco Skinny Client Control Protocol terminal control protocol messages
Routing	RIP	Routing Information Protocol distance-vector routing protocol messages
Social Media	GTalk	Google Talk or Google Chat instant text messages
	IRC	Internet Relay Chat protocol text messages
	NNTP	Network News Transfer Protocol for Usenet news articles
VoIP	RTP	Real-time Transport Protocol audio and video streaming protocol messages
	SIP	Session Initiation Protocol text-based messaging application control protocol messages
	H323	ITU-T Q.931 based protocol for Voice over IP (VoIP)
Web	HTTP/HTTP2	HyperText Transfer Protocol 1.x/2.x messages
	HTTPS/TLS	HyperText Transfer Protocol Secure over Transport Layer Security or Secure Sockets Layer encrypted messages

Note: When selective network data collection is set up to collect metadata only for the HTTPS/TLS protocol, it truncates the traffic after the SSL/TLS handshake to ensure that all the certificates and other details are retained for analysis purposes while the encrypted packet payloads are dropped.

(Optional) Configure a Decoder to Write Standard pcap-formatted Files

To provide a more open database format, the Network Decoder can write standard pcap-formatted files.


You can enable pcapng-formatted database files with the configuration node:

`/database/config/packet.file.type = 'pcapng' or 'netwitness'`

Note: This capability is enabled by default if you install 11.2 or later. If you upgrade from a previous version to 11.2, you must enable pcapng-formatted database files manually, which can result in an approximate 4% decrease in disk space (as the pcapng files require more space than the NetWitness nwdb files).

To enable writing standard pcap-formatted files:

1. Go to  (Admin) > **Services**, select a Network Decoder service, and then select  > **View > Explore**.
2. Go to **database > config**.
3. In **packet.file.type**, the default is **pcapng**.



Setting	Value
/database/config	Decoder
hash.algorithm	none
hash.databases	session,meta,packet
hash.dir	
manifest.dir	
meta.compression	none
meta.compression.level	0
meta.dir	/var/netwitness/decoder/metadb=46.63 GB
meta.dir.cold	
meta.dir.warm	
meta.file.size	auto
meta.files	auto
meta.free.space.min	1 GB
meta.index.fidelity	4
meta.integrity.flush	sync
meta.write.block.size	64 KB
packet.compression	none
packet.compression.level	0
packet.dir	/var/netwitness/decoder/packetdb=46.63 GB
packet.dir.cold	
packet.dir.warm	
packet.file.size	auto
packet.file.type	pcapng
packet.files	auto
packet.free.space.min	3 GB
packet.index.fidelity	1
packet.integrity.flush	sync
packet.write.block.size	64 KB
session.dir	/var/netwitness/decoder/sessiondb=46.63 GB

4. To change the packet file type to NetWitness formatting, type **netwitness** and press **Enter**. This change will take effect immediately on the next packet file that is created.

Note: In the pcapng database file format, the data is in clear text, and is not obfuscated by our proprietary format, which can improve security.

Caution: Please do not touch any files in the packet database directories! You must not read or edit any `pcapng` file in the packet database directories, as they are always in use while Decoder is running. Decoder always expects full and exclusive access to those files, and other processes reading those files prevent normal Decoder operation. The proper way to access the `pcapng` files is to set up a cold storage directory. This allows Decoder to copy `pcapng` files to the cold storage directory before deletion. At that point, you are responsible for managing the `pcapng` files, including making sure that the cold storage volume never fills up. Keep in mind that copying the `pcapng` files to cold storage requires a non-trivial amount of I/O and could interfere with packet capture. Cold storage for `pcapng` is not supported at 10G speeds.

(Optional) Multiple Adapter Packet Capture

Beginning with 11.5, the Network Decoder can capture from multiple interfaces simultaneously. This functionality allows Network Decoders to capture from multiple physical Network Interface Cards (NICs), or multiple ports on multiple interfaces, while leveraging the same network rules, application rules, and parsers for each NIC. The benefit of capturing from multiple physical NICs is that this method multi-threads the capture process for each adapter interface, which also multi-threads network rule evaluation.

For example, prior to 11.5, defining `capture.selected=PFRINGZC, em3` along with `capture.device.params=device=zc:em3, zc:em4` allocated a single thread for the Network Decoder to capture any traffic that came into either `em3` or `em4` interfaces. This also meant that all network rules were evaluated in a single thread.

However, in 11.5, configuring `capture.selected=PFRINGZC, em3; PFRINGZC, em4` allocates two threads for the Network Decoder to capture the traffic collected on either `em3` or `em4`, which provides more resources to the capture pipeline.


In this topic, the term "adapter" refers to a capture device within the Decoder. An adapter's name consists of the device (for example, the software API that the Decoder uses to interact with an adapter), and an interface (the physical device). A device can have multiple interfaces, and an interface can be compatible with multiple devices.

Configure Multiple Adapter Packet Capture

Note: Changes to capture device configuration do not take effect until service restart.

You can define multiple adapter packet capture for Decoders in the config setting `/decoder/config/capture.selected`. It accepts a semi-colon separated list of adapters, for example: `bpf, en5 (bpf); null_device, Null Capture`. This setting can be edited while capture is running, however, the changes do not take effect until the service has been restarted.

You can use the `/decoder?msg=select` command to select one or more adapters after capture has been stopped, for example: `/decoder?msg=select&adapter=1,3,5`



Note: As of 11.5, this configuration is not yet available in the decoder > config page. Use the decoder > explore page or the REST API. You can the Explore view by selecting a Decoder in the user interface and going to  (actions)> **View > Explore**. You can access the REST API by opening a browser and specifying the IP address of the host, for example, `https://<decoder-ip-address>:50106`.

Per Interface Configuration

The `/device/config/device.capture.params` configuration is a global configuration setting, and is applied to every interface when capture is started.

There is also a configuration setting for individual interfaces, which is `/decoder/devices/<devicename>/<interfacename>/config`. Currently, the only option available in this configuration setting is `capture.params`. Any value in this setting will override the global configuration setting.

Stat Nodes

As with configuration, there are per-interface stats and aggregate stats (similar to the global configuration). When there is a single device configured, these stat nodes function as in previous versions. When there is more than one device configured, the stat settings `/decoder/stats/capture.device` and `/decoder/stats/capture.selected` read multi.

Per-Interface Stats

The stats for each interface can be found in:
`/decoder/devices/<devicename>/<interfacename>/stats`

The captured stats are:

- `capture.avg.size`
- `capture.dropped`
- `capture.filtered`
- `capture.header.bytes`
- `capture.kept`
- `capture.payload.bytes`
- `capture.received`
- `capture.total.bytes`
- `pool.packet.captured`

Aggregate Stats

The capture stats under `/decoder/stats` represent an aggregate of the individual interface stats.

In the following table, for example, `capture.avg.size` is the maximum average size among all the interfaces, where `capture.dropped` is the total captures that were dropped for all the interfaces.

<code>capture.rate</code>	SUM
<code>capture.avg.size</code>	MAX
<code>capture.dropped</code>	SUM
<code>capture.filtered</code>	SUM
<code>capture.header.bytes</code>	SUM
<code>capture.kept</code>	SUM
<code>capture.payload.bytes</code>	SUM
<code>capture.received</code>	SUM
<code>capture.total.bytes</code>	SUM
<code>pool.packet.captured</code>	MIN

Pool Stats

All the network packet and session pool pages are tracked in `/decoder/stats/pool/`.

The stats under `/decoder/stats/pool` represents the total number of capture pages configured in Decoder. These include packet pool and session pool pages.

capture.port Meta Key

Each session is given a `capture.port` meta key which identifies the adapter from which the packets in the session were captured. This meta key can optionally be turned off with the `/decoder/config/captureport.meta.enabled` option. This meta key is on by default for Network Decoders, and is off by default for Log Decoders.

IMPORTANT: The `capture.port` meta key does not work with the `packet_mmap_, ALL` adapter. You must specify individual ports while using `capture.port`. The `packet_mmap_, ALL` adapter is capable of capturing across all types of network interfaces at the same time. For more information, see [\(Optional\) Configure a Decoder to Capture Data Across All Types of Network Interfaces](#).

Note: If you turn the `capture.port` meta key off, you must perform a capture restart for the change to take affect since this is a configuration change.

Special Devices

The `nwimport` device is for handling PCAP imports while simultaneously capturing network data. For more information on PCAP imports, see [Upload a Packet Capture File](#).

`nwimport`: This device is not selectable but does have stat and config nodes. The config node for `nwimport` is ignored.

Packet arrived out of order Message

This message means that the Decoder is seeing old packets. This can be a side-effect of running import while capture is running. This message is usually suppressed at the beginning of import and then unsuppressed at the end. However, packets can still be in the state of being processed even after the upload of the imported file has completed. This means that this warning can be unsuppressed before the processing of the imported packets is completed, which results in displaying this message.

For information about updating Decoder configurations, see the *RESTful API User Guide for RSA NetWitness Platform*. For information about stopping and restarting capture, see [Configure Capture Settings](#). Go to the [Master Table of Contents](#) to find all RSA NetWitness Platform 11.x documents.

Multi-Assembler Modes

The process that reassembles captured packets into streams is known as the Assembler. You can customize the Assembler operation to suit the packet capture configuration. It works in two modes (`on` and `off`) that are configured using the `/decoder/config/assembler.threading.enabled` variable. The value of `assembler.threading.enabled` only takes effect if multi-adapter packet capture is enabled.

Note: If you change the value of `assembler.threading.enabled`, you must perform a capture restart.

OFF Mode: If assembler threading is `off`, then one Assembler is used to assemble all packets captured from all interfaces. Packets from a stream may appear on any capture interface and the single assembler will put them back together.

- Assembler threading must be `off` to allow assembler to reconstruct streams that are capture across multiple capture interfaces.
- When assembler threading is `off`, the single Assembler instance may limit the maximum packet rate that the Decoder can process.

ON Mode: If assembler threading is `on`, there will be a separate Assembler for each capture interface. Each Assembler will only reassemble the packets it receives from it's capture interface.

- This mode of operation enables higher throughput because each Assembler instance operates on a dedicated processor.
- Each logical stream of packets must be ingested by one interface only. For example, the streams on each capture interface could come from different networks.

For compatibility with previous releases, the Multi Assembler mode is `off` by default.

When importing packets, the imported packets are fed to one of the existing assemblers, or the single Assembler instance when Assembler threading is `off`. Importing packets does not create a separate Assembler process.

(Optional) Data Plane Development Kit Packet Capture

The NetWitness Platform Decoder supports capturing packets through the Data Plane Development Kit, also known as DPDK. DPDK is a set of libraries and drivers that support applications that work on directly on packets. DPDK's web site, www.dpdk.org, provides more information.

Advantages

The Decoder can use the DPDK to ingest packets from Ethernet interfaces. It is an alternative to existing packet capture options built into the Decoder. DPDK offers these advantages:

- It is entirely open source, with no binary-only driver components, unlike PF_RING.
- It is directly supported by multiple vendors, including our main vendor, Intel.
- Since the kernel-level interface is in the mainline Linux kernel, there are no separate driver packages to compile and distribute.
- DPDK devices are completely separate from the Linux network stack, which avoids unintentional interaction between Linux networking and packet capture.
- DPDK is capable of capturing at greater than 10G, unlike the `packet_mmap`.

How it Works

DPDK completely detaches the Ethernet interface from the Linux kernel. It does this by doing the following:

- It uses a Linux User I/O driver mechanism to expose the raw PCI-Express device to normal Linux programs. This means that it can take the raw memory space mapped to the PCI device and expose it, so that a normal Linux program can directly manipulate the hardware. There are a couple of drivers built into the Linux mainline kernel that can accomplish this. Decoder is designed to use `vfio_pci_generic`.

Note: Only one driver can be attached to an Ethernet interface at a time. When `vfio_pci_generic` is attached to an interface, the Linux network driver for that interface is detached.

- The DPDK provides its own user-space implementation of the hardware drivers for a wide variety of Ethernet cards. When a DPDK application like Decoder is loaded, it can attach directly to the PCI device and manipulate it using DPDK-specific library functions. When a device is attached to DPDK, the Decoder can see it and automatically sets the device up with the DPDK Capture Device option.

Migrate PF_RING Devices to DPDK

For configurations using multiple adapter packet capture (available in 11.5 and later), you must use the manual procedure outlined in [Manually Move a Capture Interface to DPDK](#). For example, the migration (`dpdk migrate`) does not support a Decoder when capture is configured with:

```
capture.interface=PFRINGZC,em3; PFRINGZC,em4;packet_mmap,em2.
```

For configurations prior to 11.5, the procedure that uses the `devices` parameter to capture from multiple PF_RING-supported interfaces (for example, `ixgbe` and `i40e`) will be converted to using DPDK, along with the multiple adapter packet capture, so that one capture thread is used per interface.

For example, the migration will work if configured as follows:


- For multi-interface capture: `capture.interface=PFRINGZC,em3` along with `capture.device.params=device=zc:em3,zc:em4`
- For single interface capture: `capture.interface=PFRINGZC,em3`

Note: `dpdk migrate` only functions with PFRINGZC, not mmap interfaces.

If you are using PF_RING to capture packets on your Decoder, you can use the Decoder's `dpdk migrate` command to migrate your existing PF_RING capture devices to DPDK. This command uses your currently-selected capture device configuration to determine which network interfaces are moved to DPDK control.

1. Go to  (Admin) > Services and select a Decoder.
2. Click  > **View** > **Explore**, right-click **decoder** and select **Properties**.
3. From the command drop-down list, select the `dpdk` command, type the parameter `migrate`, and click **Send**.

- In the command output window, the changes that will be made on the Decoder to perform the migration are displayed. If everything looks correct for the migration, add the parameter `commit=1` to the command after `migrate` to actually make the changes on the Decoder.



- At the reboot prompt, reboot the host to make the driver loading changes.
- (Optional) Run the `/decoder/devices/<devicename>/msg=prune` command to remove configuration options for devices and interfaces that no longer exist after the migration.

Note: The `prune` command removes only the undetected interface folders. Also, it ignores the `interfaces`, `stats`, `nwimport`, and `flow_events` folders available in `/decoder/devices/`.

Manually Move a Capture Interface to DPDK

- Run the new host service command `/appliance/nicToDp`, that converts a named network interface and sends it to DPDK.
For example, if you provide `nicToDp` an interface name like `p1p1` it detaches it and sets it up for DPDK to use. It does the following:

- a. Identifies the PCI address of the named interface.

Note: From this point on, the Ethernet interface is identified by its PCI address alone, which is a physical characteristic based on where the Ethernet interface is installed in the computer.

- b. Creates configuration files that describe which PCI addresses DPDK is allowed to use.
 - c. Sets up a systemd unit that binds the configured PCI addresses to DPDK instead of to the default Linux drivers on the next system reboot.
2. Reboot the Decoder host. The reboot accomplishes the following:
 - a. Ensures that live capture on the interfaces stops before the network driver is unloaded.
 - b. Binds the `vfio_pci_generic` driver to the configured interfaces.
 - c. Pre-allocates some Huge Page memory for each interface, as required by the DPDK libraries.
3. After the reboot, the Decoder does not start capturing, since the network interface it had been using to capture has been moved to DPDK. Change the Decoder's capture configuration to select the DPDK device that is now available. For information about how to change the capture configuration, see [Configure Capture Settings](#).
4. Restart the service.

Utilizing Receive Side Scaling with DPDK

Receive Side Scaling (RSS) efficiently distributes the network receive processing across multiple CPUs in a multi-processor system. RSS ensures that the processing associated with a given connection stays on the assigned CPU.

IMPORTANT: Receive Side Scaling is supported on DPDK devices using the ixgbe or i40e device drivers only.

Advantages

When RSS is enabled, the DPDK capture devices split the incoming packet stream into multiple queues providing following benefits:

- Each queue is processed with a dedicated capture worker thread. As there are more threads processing the capture work, it allows more CPU time available for Berkeley Packet Filters (BPF) rule and network rule evaluation.
- Each capture worker thread can feed multiple assemble worker threads simultaneously which allows more CPU time available for assembly.
- Each CPU thread processes only a fraction of the packet stream. However, the total number of packets processed from all queues working in parallel increases the overall output (packets processed per second) of Decoder.

Configuring Receive Side Scaling

To enable RSS in a DPDK controlled NIC, set the `/decoder/devices/DPDK/(PCI_ID)/config/rss` value to a value greater than 1. The typical values of RSS must be a small power of 2, for example, 4 or 8. This value determines the number of CPU threads to be allotted for packet processing during the next capture. Similarly, if assembler threading is enabled, a matching number of assembler threads will be created to process the packets on each incoming queue.

For example, the following sample configuration will utilize a total of eight threads (four capture threads + four assembler threads) for multi-adapter packet capture with DPDK and multi-assembler threading turned on. This is the maximum recommended value for an eight core Decoder CPU.

```
assembler.threading.enabled set to on or true
/decoder/devices/DPDK/0000:01:00.1/config/rss=2
/decoder/devices/DPDK/0000:01:00.0/config/rss=2
```

To minimize packet drops, it is recommended that you keep the total number of capture and assembler threads less than the number of cores per CPU.

How it Works

When RSS is enabled, Decoder utilizes the RSS hardware acceleration on the capture NIC to apply a symmetric RSS hash function. The hash function sends packets from both directions of a given stream to the same queue. It enables the packets received on that queue to be processed by a dedicated assembler thread. This eliminates the need for Decoder to merge streams from different RSS queues and reassemble the streams. Hence, it is recommended that you enable multi-threaded assembly when using RSS. To enable, set `/decoder/config/assembler.threading.enabled` to `true` or `on`.

(Optional) Preserve VLAN Tags When Using the Packet MMAP Capture Interface

When capturing traffic containing VLAN tags, you may need to configure the Packet MMAP capture interface to preserve the VLAN tags in the packets (VLAN fixup). By default, the network capture hardware removes the tags. Performing this procedure preserves the tags in the packets, and the tag values are parsed into VLAN meta data for further analysis.

There are two mechanisms for enabling the VLAN fixup.

- Option 1: Set `vlan-fix=true` within `capture.device.params`. This option performs the VLAN fixup on all traffic entering the Decoder. This option is appropriate in most cases, since it is assumed that all the traffic will be VLAN tagged. This mechanism works on either single-interface mode, or on all-interfaces mode. This option overrides the VLAN fixup settings on individual interfaces; even interfaces that are not configured to do VLAN fixup will have the feature enabled.
- Option 2: Use the `interfaces` parameter within `capture.device.params` on a per-device basis. The `interfaces` parameter accepts a comma-separated list of interface names on which to capture packets. By adding `:vlan` to an interface name, you can enable the VLAN fixup on individual interfaces. If the interface does not have the `:vlan` suffix added, then it will not perform the VLAN fixup.

After editing this parameter, you must restart capture on the Decoder in order for changes to `capture.device.params` to take effect.

These are `vlan` examples of both options. If you need to pass multiple settings for `capture.device.params`, use the following syntax. Notice that quotes are needed to delineate whitespace. For more information about how to use quotes for whitespace, see the "Connecting to a Service" topic in the *NwConsole User Guide for RSA NetWitness Platform*.

```
name1="value1" name2="value2".
```

Parameter	Value	Effect
<code>capture.device.params</code>	<code>vlan-fix=true</code>	VLAN fixup always performed on all interfaces. The default value is <code>vlan-fix=false</code> .
<code>capture.device.params</code>	<code>interfaces=eth0:vlan,eth1</code>	VLAN fixup performed on traffic capture on <code>eth0</code> interface only
<code>capture.device.params</code>	<code>interfaces=eth0:vlan,eth1</code> <code>vlan-fix=true</code>	VLAN fixup always performed because the <code>vlan-fix</code> setting overrides the interfaces setting.

To configure the `packet_mmap_adapter` to preserve the VLAN tags in packets:

1. Go to  (Admin) , select the Decoder service and  > **View** > **Config**.

- In the **Decoder Configuration** panel, set **Capture Interface Selected** to the `packet_mmap_`, ALL adapter.

The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes tabs for Investigate, Respond, Users, Hosts, Files, Dashboard, and Reports. Below this is a secondary navigation bar with tabs for Hosts, Services, Event Sources, Endpoint Sources, Health & Wellness, System, and Security. The main content area is divided into two panels: System Configuration and Decoder Configuration.

System Configuration

Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☒
SSL Port	56004
Stat Update Interval	1000
Threads	20

Decoder Configuration

Name	Config Value
Adapter	Berkeley Packet Filter
Capture Interface Selected	packet_mmap_eth0 (bpf)
Cache	Cache Directory
Cache Size	Cache Size
Capture Settings	Assembler Maximum Size
	Assembler Minimum Size
	Assembler Session Flush
	Assembler Session Pool
	Assembler Timeout Packets
	Assembler Timeout Session

Parsers Configuration

Specify if relevant meta data is generated to disk (Enabled), generated only in memory for other Decoder content use (Transient), or not generated at all (Disabled).

Name	Config Value
ALERTS	Enabled
DHCP	Enabled
DNS	Enabled
Entropy	Disabled
FeedParser	Enabled
FTP	Enabled
GeoIP2	(Mixed)
GTalk	Enabled
H323	Enabled
HTTP	Enabled
HTTP2	Enabled
HTTPS	Enabled
IRC	Enabled
JSNormalize	Enabled
MAIL	Enabled
NETBIOS	Enabled
NETWORK	Enabled
NFS	Enabled
NNTP	Enabled
OPENAPPID	Enabled
PGP	Enabled

At the bottom of the Decoder Configuration panel is an **Apply** button.

- To go to the **Explore** view, click **Config** in the toolbar and select **Explore** in the drop-down list.

4. In the Services Explore view select **decoder > config**.

Path	Value
/decoder/config	Decoder
adapter.meta.enabled	on
aggregate.precache	on
assembler.parse.valve	false
assembler.pool.ratio	70
assembler.session.flush	1
assembler.session.pool	2000
assembler.size.max	32 MB
assembler.size.min	0
assembler.tcp.close	false
assembler.timeout.packet	60
assembler.timeout.session	60
assembler.voting.weights	first=1 size=1 port=1 octet=1 routable=1
capture.autostart	on
capture.buffer.size	8 MB
capture.device.params	
capture.selected	packet_mmap_eth0 (bpf)
export.cache.expire	60
export.packet.enabled	no
export.packet.local.path	/var/netwitness/decoder/export/packet
export.packet.max	0
export.packet.remote.path	/var/netwitness/decoder/export/packet
export.packet.size.max	1 GB
export.session.enabled	no
export.session.format	avro

5. Click in the values column next to `capture.device.params`, and do one of the following:

- To preserve VLAN tags on an interface in the interfaces list, add **:vlan** after the interface name and press **Enter**.


For example, this specifies that VLAN tags are preserved on `em1`, but not on `em2` and `em4`:

`interfaces=em1:vlan,em2,em4`

Parameter	Value
adapter.meta.enabled	on
aggregate.precache	on
assembler.parse.valve	false
assembler.pool.ratio	70
assembler.session.flush	1
assembler.session.pool	2000
assembler.size.max	32 MB
assembler.size.min	0
assembler.tcp.close	false
assembler.timeout.packet	60
assembler.timeout.session	60
assembler.voting.weights	first=1 size=1 port=1 octet=1 routable=1
capture.autostart	on
capture.buffer.size	8 MB
capture.device.params	interfaces=em1:vlan,em2,em4
capture.selected	packet_mmap_eth0 (bpf)
export.cache.expire	60
export.packet.enabled	no
export.packet.local.path	/var/netwitness/decoder/export/packet
export.packet.max	0
export.packet.remote.path	/var/netwitness/decoder/export/packet
export.packet.size.max	1 GB
export.session.enabled	no
export.session.format	avro

The change goes into effect after restarting capture; only traffic on em1 has the VLAN tags preserved.

- To preserve VLAN tags on all interfaces, enter the following and press **Enter**:
vlan-fix=true



VLAN tags are preserved on all capture interfaces.

6. Select Start Capture (**Start Capture**). The change takes affect after capture is restarted.


(Optional) Process Raw Syslog Data without Priority Field

You have the option to process raw syslog data that does not contain a valid priority (PRI) field.

To configure a Log Decoder to process syslog without a Priority field:

1. Go to (Admin) > **Services**, select the Log Decoder service and > **View** > **System**.
2. Select Stop Capture (**Stop Capture**).
3. From the drop-down menu, where System is shown, select **Explore**.
4. Select **decoder** > **config**.
5. In the `capture.device.params` field, add the following text, and then click **Enter** to save the changes:

```
requirePri=false
```



- From the drop-down menu, where Explore is shown, select **System**.
- Select Start Capture (**Start Capture**). The change takes affect after capture is restarted.

(Optional) Configure Decoder to Support OpenAppID

The NetWitness Platform Decoder can identify applications using the OpenAppID detectors. OpenAppID from Cisco is an application-layer network security plug-in for Snort (an open source network intrusion detection system). It is a set of open source Lua libraries (detectors) that identifies applications in the network traffic.

When Network Decoder detects an application with an OpenAppID detector, the `appid` meta is created. The value of the `appid` meta helps in identifying the application as it is the application's entry in the `appMapping.data` file. The `appMapping.data` file is part of the OpenAppID distribution and contains metadata of the application detectors.

Note: By default, all detectors are disabled to avoid any negative impact on the performance.

To configure the OpenAppID detectors:

IMPORTANT: Ensure that you enable only the required OpenAppID detectors. Enabling all the OpenAppID detectors can reduce the performance. Before you enable any detector, understand your application detection requirements and enable only the specific detectors. For example, enabling the `service_*` detector or any other detector that attempts to match short generic patterns can cause severe impact on the performance.

1. Download the OpenAppID source package (snort-openappid.tar.gz) from the [Snort Downloads \(https://www.snort.org/downloads\)](https://www.snort.org/downloads) page.
2. Upload the OpenAppID source package to the network decoder (`http://<networkdecoder>:50104/decoder/parsers/upload`), which can be accessed using the REST interface.
Once uploaded, decoder will extract the package to an appropriate folder. The default folder is `/etc/netwitness/ng/odp`.
3. To enable detectors, modify the `/decoder/parsers/config/openappid.enabled` file with a combination of detector filenames or wildcards, separated by commas. For example, `client_*`, `payload_slashdot.lua`, `payload_reddit.lua`.
The following table describes the detector group filenames and their functions. You can extract the complete set of detectors from the OpenAppID source package (snort-openappid.tar.gz) or it can be viewed on the Decoder file system at `/etc/netwitness/ng/odp/lua`.

Detector Group Filename	Description
<code>client_*</code>	Detectors for client applications (for example, Facebook, MacAppStore).
<code>content_group*</code>	Detectors for a specific type of network attribute such as protocol, port, or service (for example, <code>port_services</code> which contains many services identified by their common ports).
<code>payload_*</code>	Detectors for the host based traffic (for example, Hulu, Reddit).
<code>service_*</code>	Detectors for service applications (for example, Minecraft, VNC).
<code>ssl_host_group_*</code>	Detectors for SSL hosts logically grouped together (for example, 334 which contains many ecommerce websites).

4. Ensure that you specify the detector filenames correctly. If you specify an incorrect filename, the detector will not load. The following table lists some good and bad file naming conventions and their status on subsequent loading.

Detector Group Filename	Status
<code>client_9P.lua</code>	loaded
<code>client_9p.lua</code>	not loaded (Linux only)
<code>client_9P</code>	not loaded
<code>client_9p*</code>	loaded
<code>client_</code>	loaded

Detector Group Filename	Status
9P.*	
lient_9P.lua	not loaded
lient_9P.*	not loaded
*lient_9P.lua	loaded

- Once the required detectors are enabled, reload the parser to load the detectors. The following example message will confirm that the detectors are enabled successfully.

```
2020-May-14 17:03:37 [OpenAppID] 59 detectors loaded out of 59
```

Enable and Disable Parsers and Log Parsers

Administrators can see which parsers have been downloaded from Live and deployed on a Decoder or Log Decoder, see which of these have been enabled, and enable or disable parsers and log parsers.

The following figure illustrates commonly used settings on a Decoder. For a quick basic setup with only the required steps, see [Decoder and Log Decoder Quick Setup](#).



You should only download and deploy the parsers you need for the following reasons:

- There is an impact on performance as you increase the number of deployed parsers.
- The more parsers you deploy, the more meta data created, which impacts data retention.
- Not having extra (unnecessary) log parsers deployed reduces the potential for mis-identification of messages.

The Parsers Configuration panel provides a way to select parsers to use on the Decoder. Within some parsers, you can also configure the metadata that the parser creates. These are the options in the Parsers Configuration panel.

Option	Description
Enable All Disable All	These options provide a way to quickly select either all parsers or no parsers.
Name	The names of parsers available to the Decoder. A plus sign indicates that the metadata generated by the parser is configurable. Clicking the plus sign displays the metadata that the parser can create.
Config Value	A drop-down list changes the setting for the parser or metadata to Enabled, Disabled, or Transient. • When Enabled, the Decoder is using the parser to filter traffic. • When Transient, the Decoder is using the parser to filter traffic, and the generated metadata is not stored on disk. The transient metadata is available in memory to additional content (that is, parsers, feeds, and application rules) on that Decoder. This helps administrators to protect certain data and is usually done as part of a data privacy plan (see the <i>Data Privacy Management Guide</i>). • When Disabled, the Decoder is not using the parser. If the generated metadata for the parser is configurable, clicking the plus sign to expand the parser displays configurable meta keys and the same drop-down list selects the meta key the parser will create.

Note: For a Log Decoder, you must have previously deployed log parsers from Live. See the **Find and Deploy Live Resources** topic in the *Live Services Management* Guide for details. Go to the [Master Table of Contents](#) to find all RSA NetWitness Platform 11.x documents.

To enable or disable a parser, or to view the status for each parser:

1. Go to  (Admin) > **Services**.
2. In the **Administration Services** view, select a Log Decoder or a Decoder, and  >**View** > **Config**.
3. In the **Parsers Configuration** panel, look for the Decoder parser or the Log Decoder event source parser.

Parsers Configuration		Enable All	Disable All
Specify if relevant meta data is generated to disk (Enabled), generated only in memory for other Decoder content use (Transient), or not generated at all (Disabled).			
Name	Config Value		
 ALERTS	Enabled		
alert	Enabled Disabled Transient Enabled		
 DHCP			
 DNS			
 Entropy			
FeedParser	Enabled		
 FTP	Enabled		
 GeoIP2	(Mixed)		
 GTalk	Enabled		
 H323	Enabled		
 HTTP	Enabled		
 HTTPS	Enabled		
 IRC	Enabled		
 MAIL	Enabled		
 NETBIOS	Enabled		
 NETWORK	Enabled		
NFS	Enabled		
 NNTP	Enabled		
 PGP	Enabled		

4. In the **Config Value** column, note the current status for your parser.
 You can update the status of any individual parser by selecting its **Config Value** and selecting **Disabled**, **Transient**, or **Enabled** from the drop-down menu. Alternatively, you can select **Enable All** or **Disable All** to update the status for all of your log parsers at once.
5. Click **Apply**.

When you click **Apply**, note that all parsers are reloaded into NetWitness Platform. The status for each parser is updated, based on your selections.

Start and Stop Data Capture

When a Decoder starts up, it automatically begins aggregating data if **Capture Autostart** is enabled. When autostart is not enabled, you can start and stop data capture manually.

Note: The Capture Configuration Settings in the Service Config view for a Decoder determine whether Capture Autostart is enabled.

The following figure illustrates commonly used settings on a Decoder. For a quick basic setup with only the required steps, see [Decoder and Log Decoder Quick Setup](#). You may want to stop and start capture at other times, for example, before you shut down the service.



To start and stop capture:

1. Go to  (Admin) > **Services**.
2. Select a Decoder or Log Decoder service, and select  > **View** > **System**.
3. In the toolbar, click **Start Capture**.

If the service is a Decoder, it begins capturing packets. If the service is a Log Decoder, it begins capturing logs.

When packet or log capture is in progress, the option in the toolbar changes to **Stop Capture**, and the option to upload a file is unavailable.

4. Whenever you want to discontinue traffic capture on a Decoder, click **Stop Capture**.


Packet or log capture ceases, and the option to upload a file to the service is again available.

Note: When you stop the Log Decoder service while capture is running, all events currently in Log Decoder memory will be processed and persisted. Should an issue arise where it is necessary to quickly shutdown the service, use the Services Explore view to stop capture (`/decoder stop`), passing the parameters `flush=false` before stopping the Log Decoder service. For further information, see the "Services Explore View" in the *Host and Services Getting Started Guide*.

Configure Decoder Rules

This topic provides procedures for creating and managing rules for Decoder or Log Decoder traffic capture in the **Services Config view > Rules** tabs. [Services Config View - Rules Tabs](#) provides details about the Rules tab options.

The following figure illustrates commonly used settings on a Decoder. For a quick basic setup with only the required steps, see [Decoder and Log Decoder Quick Setup](#).



Capture rules can add alerts or contextual information to sessions or logs. They can also define which data is filtered out by a Decoder or Log Decoder. Rules are created for specific metadata patterns, which result in predefined actions when matches are found. For example, to keep all traffic that fits certain criteria, but discard all other traffic, you can create a rule to perform the necessary actions. When applied, rules affect both packet capture file importing, as well as live network capture.

[Rule and Query Guidelines](#) provides guidelines that all queries and rule conditions in NetWitness Platform Core Services must follow.

By default, no rules are defined when you first install NetWitness Platform. Until rules are specified, the packets are not filtered. You can deploy the latest rules from Live. You can define three types of rules: Network Rules, Application Rules, and Correlation Rules.

- Network rules are applied at the packet level and are made up of rule sets from Layer 2, Layer 3, and Layer 4. Multiple rules can be applied to the Decoder. Rules can be applied to multiple layers (for example, when a network rule filters out specific ports for a specific IP address). Network rules are only available on Network Decoders.
- Application rules are applied at the session level. If the first rule listed is not a match, the Decoder then attempts to match the next rule listed, until a match is found.
- Correlation rules are applied over a configurable sliding time window. When a match is found, the service creates a new super session that identifies other sessions that match the rule, then creates a session list for analysis.

The two most common uses of rules are:

- To alert, and thereby create a custom alert meta value, when certain conditions are found.
- To filter out certain types of traffic that do not add value to the analysis of the data.

Groups of capture rules form rule sets, which you can import and export. This feature enables use of multiple rule sets for various scenarios. You can import the exported rule set, in the form of an .nwr file, to other NetWitness Platform services, simplifying the deployment and configuration of multiple services.

Rule Processing

These are the principles governing capture rule processing:

- Multiple rules can be applied to the Decoder.
- Capture rules are executed one after the other, in sequence.
- Rule processing stops when all rules are processed or after a rule configured to stop rule processing is matched.
- A default rule can be used to either include or exclude all traffic not otherwise selected by a rule. A default rule, if used, must always be placed at the bottom of the rule list. Otherwise, rule processing stops as soon as the default rule is evaluated since, by definition, all traffic is selected by the default rule.
- When rule processing stops, the session is saved using the configured session options and debug options.

Rule and Query Guidelines

All queries and rule conditions in RSA NetWitness Core services must follow these guidelines:

All string literals and time stamps must be quoted. Do not quote numbers, MAC, or IP addresses.

- `extension = 'torrent'`
- `time='2015-jan-01 00:00:00'`
- `service=80`
- `ip.src = 192.168.0.1`

Note: The space on the right and the left of an operator is optional. For example, you can type a rule as `service=80` or `service = 80`.

Rule Examples

The following table shows examples of rule conditions. You can use rule conditions for log retention collections in an Archiver and for application, network, and correlation rules on a Decoder, Log Decoder, or Concentrator. Rule conditions are also used in all `WHERE` clauses in all Core database queries.

For detailed information on rule syntax in NetWitness Platform, see "WHERE Clauses" in the "Queries" section of the *Core Database Tuning Guide*.

Rule Name	Condition
ComplianceDevices	<code>device.group='PCI Devices' device.group='HIPPA Devices'</code>
HighValueWindows	<code>device.group='Windows Compliance'</code>
MediumValueWindows	<code>device.type='winevent_nic' && msg.id='security_4624_security'</code>

Rule Name	Condition
LowValueWinLogs	<code>device.type='winevent_nic' && msg.id='security_4648_security'</code>
LowValueProxyLogs	<code>device.class='proxy' && msg.id='antivirus_license_expired'</code>
GeneralWindows	<code>device.type='winevent_nic'</code>

Invalid Rules

NetWitness Platform uses a rule parser that strictly defines valid syntax for rules and queries. When a Core service encounters invalid syntax, it writes a warning in the NetWitness Platform logs indicating the error.

Note: NetWitness Platform 11.x does not support parsing of legacy syntax rules (as version 10.6 did). After you update to NetWitness Platform 11.x, rules with invalid syntax are highlighted in the user interface, and no rules will be applied until the invalid rules are corrected. The Rule Editor provides additional tooltips. After you fix the rules, the highlights disappear. See [Fix Rules with Invalid Syntax](#).

The `/decoder/config/rules/rule.errors` and `/concentrator/config/rules/rule.errors` stats, contain the count of rules with errors. If `rule.errors` is nonzero, NetWitness Platform generates a Health and Wellness alert to indicate that you need to fix the rules.

General Syntax Guidelines

- All text values must quote literal values. Example: `user = 'user1'`
- Quotes can use single or double quotes; but they must match. (You cannot start with a single quote and finish with a double quote.)
- If the literal value has a quote, you can escape it using double quotes. Example: `user = "User's"`

The following are valid syntax rules:

- All time values should use quotes for dates in this form: `time = 'YYYY-MM-DD HH:MM:SS'`
- All time values that are the number of seconds since EPOCH (Jan 1, 1970), should not be quoted. Example: `time = 1448034064`
- Everything else is unquoted: IP addresses, MAC addresses, numerics, and so on. Example: `service = 80 && ip.src = 192.168.1.1/16`

Capture Rule Syntax

Capture rules compare fields to values or to other fields. This is an example of a simple expression with a meta key on the left side of the operator and a value on the right side.

```
ip.dst=192.168.1.1
```

The syntax allows a meta key on the right side of the operator in Decoders and Log Decoders for application and network rules. Meta key comparison does not apply in the `where` clause in queries. This is an example of a simple expression with a meta key on the left side of the operator and a meta key on the right side.

```
ip.src=ip.dst
```

Rules that include a meta key comparison support renamed meta keys; if a rule queries a meta key that has been renamed, the rule is parsed for the renamed meta key. For example, if the meta key `ip_dst` is used in a rule, it is transparently mapped to the renamed meta key: `ip.dst`. Existing rules that include original keys will trigger alerts that include data for the renamed meta key.

This is an example of a rule that finds packets having the same `ip.src` address and `ip.dst` address on a Decoder, and generates an alert on the Concentrator.

```
alert=ioc name=testRule8 rule="ip.src=ip.dst" order=38
```

This rule would generate an error because `eth.src` and `ip.src` are incompatible formats.

```
rule="eth.src=ip.src" name="testRule99" alert=ioc
```

Values can be expressed as discrete values, a range of values, an upper or lower bound, or a combination of these three. You can create a greater than or less than comparison, and test equality or inequality against a range of values or an upper/lower bound.

`key 0-5` (a range of values)

`key = 0-u` is the same as `key >= 0` (upper bound, greater than or equal to)

The following table summarizes the operators on meta keys.

Left Operand Format	Operator	Right Operand Format	Description
any	=	compatible with left operand	Equality operator. You can use values or meta keys on the right side of the equality operator.
any	!=	compatible with left operand	Inequality operator. You can use values or meta keys on the right side of the inequality operator.
any	<	compatible with left operand	Less than operator. You can use values or meta keys on the right side of this operator.

Left Operand Format	Operator	Right Operand Format	Description
any	<code><=</code>	compatible with left operand	Less than or equal to operator. You can use values or meta keys on the right side of this operator.
any	<code>></code>	compatible with left operand	Greater than operator. You can use values or meta keys on the right side of this operator.
any	<code>>=</code>	compatible with left operand	Greater than or equal to operator. You can use values or meta keys on the right side of this operator.
text	<code>contains</code>	text	Find values that contain the right operand. You can use meta keys or values on the right side of this operator.
text	<code>begins</code>	text	Find values that begin with the right operand. You can use meta keys or values on the right side of this operator.
text	<code>ends</code>	text	Find values that end with the right operand. You can use meta keys or values on the right side of this operator.
text	<code>length</code>	integer	Find strings of a certain length. You can use meta keys or values on the right side of this operator.
any	<code>count</code>	integer	Find values with a specific number of occurrences within the session. You can use meta keys or values on the right side of this operator.
any	<code>ucount and unique</code>	integer	Finds a number of uniquely occurring values. You can use meta keys or values on the right side of this operator. For example, if the results include instances of a meta key with five unique values and three of the same value, the <code>ucount</code> is six.
N/A	<code>exists</code>	any	Finds any values for the meta key. You can use meta keys or values on the right side of this operator.
N/A	<code>!exists</code>	any	Finds any sessions in which the meta key does not occur. You can use meta keys or values on the right side of this operator.
text	<code>regex</code>	text	Finds values matching a regular expression. You can use values on the right side of this operator.

The following table summarizes other syntax elements used in rules.

Syntax element	Description
*	Default rule. By using an asterisk (*) as the sole character in a rule, that rule will select all traffic.
u	Upper bound of a range of times, IP addresses, or numeric formats. For example, to select all TCP ports above 40000, the syntax would be: <code>tcp.port = 40000-u</code>
l	Lower bound of a range of times, IP addresses, or numeric values. For example, to select all TCP ports below 40000, the syntax would be: <code>tcp.port = l-40000</code>
- (dash)	Denotes a range. This is only applicable to time values, IP or MAC addresses, or numeric values. Separate the lower and upper bounds of the range with a dash (-) character. For example, to select TCP ports between 25 and 443, the syntax would be: <code>tcp.port = 25-443</code>
, (comma)	Denotes a list of ranges or values or meta keys. Single values may be used as well as any combination of ranges and upper or lower bounds. Single meta keys may be used in a list. Meta keys and literal values cannot both appear on the right-hand side of an operator. For example, the following is valid syntax: <code>tcp.port = l-10,25,110,143-225,40000-u</code>
()	Grouping operator. An expression can be enclosed in parentheses to create a new logical expression. For example, the following would select traffic on port 80 to/from 192.168.1.1 OR traffic on port 443 to/from 10.10.10.1: <code>(alias.ip=192.168.1.1 && tcp.port=80) (alias.ip=10.10.10.1 && tcp.port=443)</code>
~	Logical NOT operator, a negation of an expression.
&&	Logical AND operator, a conjunction of two expressions.
	Logical OR operator, a disjunction of two expressions.

Configure Capture Rules

The Decoder and Log Decoder rules are editable in the Services Config view. While each type of rule (network, application, and correlation) has its own tab; the functions are similar for all types of rules. You can:


- Add, edit, and delete rules
- Enable and disable rules
- Change the execution sequence of rules
- Import rules from a file
- Export rules to a file

- Push rules to another service
- Revert or apply rule changes
- Restore one of the last ten rule configurations from a snapshot

To configure rules in the Rules tabs

1. Go to  (Admin) > **Services**.
2. In the **Services** view, select a Decoder service and  > **View > Config**.
3. In the **Services Config** view, select one of the Rules tabs: Network Rules, App Rules, or Correlation Rules.

The rules list for the selected rule type is displayed.



Status	Pending	Name	Condition	Session Data	Alert	Last Updated By
☐	☒	nw110025	ip.proto=1 && (filename exists filetype exists)		alert.id	
☐	☒	nw110030	service=0 && (filename exists filetype exists)		alert.id	
☐	☒	nw60005	service = 53 && udp.dstport = I-52,54-u && strea...		alert.id	
☐	☒	nw60010	service = 23 && tcp.dstport = I-22,24-u && strea...		alert.id	
☐	☒	nw60015	service = 21 && tcp.dstport = I-20,22-u && strea...		alert.id	
☐	☒	nw60020	service = 80 && tcp.dstport = I-79,81-u && strea...		alert.id	
☐	☒	nw60025	service = 22 && tcp.dstport = I-21,23-u && strea...		alert.id	
☐	☒	nw60030	service = 25 && tcp.dstport = I-24,26-u && strea...		alert.id	
☐	☒	nw60035	service = 67 && tcp.dstport = I-66,68-u && strea...		alert.id	
☐	☒	nw60040	service = 69 && tcp.dstport = I-68,70-u && strea...		alert.id	
☒	☒	nw60045	service = 110 && tcp.dstport = I-109,111-u && ...		alert.id	
☐	☒	nw60050	service = 119 && tcp.dstport = I-118,120-u && s...		alert.id	
☐	☒	nw60075	service = 443 && tcp.dstport = I= 443,993,995 && ...		alert.id	
☐	☒	nw60080	service = 520 && tcp.dstport = I-519,521-u && s...		alert.id	
☐	☒	nw60085	service = 1433 && tcp.dstport = I-1432,1434-u && ...		alert.id	
☐	☒	nw60090	service = 1521 && tcp.dstport = I-1520,1522-u && ...		alert.id	
☐	☒	nw60095	service = 1719 && tcp.dstport = I-1718,1720-u && ...		alert.id	

Each type of rule has a list with slightly different columns and different parameters. Several basic guidelines apply to all rule management activities:

- The rules are executed in the sequence they are displayed in the list. To change the execution sequence of rules, drag and drop rules to the appropriate location in the list or use the context menu options to arrange the rules in the list.
- To select a single row, click the row.
- To select a group of adjacent rows, click the first, then shift-click the row at the end of the group.
- To select multiple non-adjacent rows, click the first, then control-click the others.
- When editing rules in the Rules tab, you must apply the configuration changes in order to activate.

- Until changes are applied, you can discard edits to the list and revert to the unedited rules.
- Once rules are applied, you can recover the last ten rules configurations using the **History** option in the **Actions** menu.

To add a rule in any Rules tab, do one of the following:

- Click  .
- Right-click a rule, and select **Insert Above** or **Insert Below** from the context menu.
The Rule Editor dialog for that type of rule is displayed.

To remove a rule:

1. From any Rules tab, select the rules to remove from the rules list.
2. Click  .
The selected rules are removed from the list, but still exist on the service.

To edit a rule

1. From any Rules tab, select the rule to edit.
2. Click  or double-click the rule row.
The Rule Editor dialog for that type of rule is displayed.

To disable a rule:

1. From any Rules tab, select the rules to disable.
2. Click  **Disable** .
The status changes to disabled in the rules list, but the rule is still enabled on the service.

To enable a rule:

1. From any Rules tab, select the rules to enable.
2. Click  **Enable** .
The status changes to enabled in the rules list, but the rule is still disabled on the service.


Import Rules from a File and Export Rules

You can import network, application, and correlation rules to a Decoder from a file that contains rules of the same type. After the rules are imported, you can edit and manage them as you would any other rules.


When you attempt to import a group of rules, NetWitness Platform Administration checks the type of rules imported. If you are successful, a message displays the number of rules imported. If the rule type differs from the active tab type, the rules are not imported. You must re-import the rules under the correct tab or select another file to import.

To import rules to a service:

1. From any Rules tab, select  **Import**  **Import** .
The Import dialog is displayed.



2. Click  .
A view of the directory structure is displayed.
3. Choose one or more NetWitness rules (.nwr) files to import, and click **Open**.
The file is added to the list in the Import dialog.



4. Click **Import**.
The rules are imported into the user interface. Imported rules have a red corner in each edited column.
5. Edit or reorder the rules if needed.
6. To save the rules to the service, click **Apply**.
The rules for the service are updated with the changes.

To export a rule to a file:

1. To export a subset of the rules, select the rules to be exported.
2. Do one of the following:
 - In the toolbar, select  **Actions** ▾ > **Export** > **Selection**. (**Export** > **All** exports all rules in the rules list even if you have a subset selected for export.)
 - Right-click the selected rules and select **Export Selection**.
A prompt for the filename is displayed.
3. Enter the filename and click **Export**.
The **.nwr** file is downloaded.


Push Rules to Other Services

You can apply (push) rules or selected rules to other services (Decoders or Log Decoders) or service groups. When you push all rules to other services, all rules on the target services are removed and replaced with all of the rules on the source service.

To push selected rules from this Decoder to other Decoders:

1. From any Rules tab, select the rules that you want to push to another Decoder.
2. Do one of the following:
 - Select  **Actions** ▾ > **Push** > **Selection**.


- Right-click the selected rules and select **Push Selected Rules**.
The Push Selected Rules dialog is displayed.



3. Select a Push Option:
 - Select **Replace All** to delete all rules on the target services and replace them with the selected rules. This is the default selection.
 - Select **Merge** to merge the selected rules with the existing rules on the target services.
4. On the **Services** tab, select the target services to receive the pushed rules, or select the groups of services from the **Groups** tab.
5. Click **Push**.
The rules are pushed to the selected services and become effective immediately.

To push all rules from this Decoder to other Decoders:

1. From any Rules tab, select  **Actions** > **Push** > **All**.
(**Push** > **All** pushes all rules in the rules list even if you have a subset selected to push.) The Push Selected Rules dialog is displayed.



- On the **Services** tab, select the target services to receive the pushed rules, or select the groups of services from the **Groups** tab.
- Click **Push**.
All rules from the target services are deleted and replaced with all of the rules from source service. The rules become effective immediately.

Change Execution Order of Rules

Capture rules are applied in the order they are displayed in the rules list. To reorder rules, use either of these methods:


- Drag and drop the rules in the appropriate location in the rules list.
- Right-click a rule to display the context menu, and use the **Cut** and **Paste** options.

Restore a Rule Snapshot from History

NetWitness Platform keeps the last ten snapshots of rules applied to a service.

To restore a rules snapshot from history:

- Select **Actions** > **History**.
A submenu of snapshots is displayed.



2. Select the snapshot time from the submenu.
The rules from the snapshot are loaded into the rules list, replacing the current set. But the current set is still in use on the service.
3. To apply the rules to the service, click **Apply**.
The rules are applied to the service.

Configure Application Rules

Application layer rules are applied at the session level. The following are sample application rules.

To truncate packets carried via Server Message Block protocol (SMB), create a rule as follows:


- Rule Name: Truncate SMB
- Condition: `service=139`
- Rule Action: Truncate All

To retain email to and from a specific e-mail address, create a rule as follows:

- Rule Name: Retain Email Tom Jones
- Condition: `email='Tom.Jones@TheShop.com'`
- Rule Action: Keep

To add or edit an application rule:

1. Go to  (Admin) > **Services**.
2. Select a Decoder or Log Decoder service and  > **View** > **Config**.
The Systems Config view for the selected service is displayed.
3. Select the **App Rules** tab.




The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. The main navigation bar has tabs for 'HOSTS', 'SERVICES', 'EVENT SOURCES', 'ENDPOINT SOURCES', 'HEALTH & WELLNESS', 'SYSTEM', and 'SECURITY'. The 'SERVICES' tab is active, and the 'NH - Decoder' service is selected. The 'Config' sub-tab is active, and the 'App Rules' tab is selected within the configuration view.

The 'App Rules' tab displays a table of application rules. The table has columns for 'Status', 'Name', 'Condition', 'Session Data', 'Alert', and 'Last Updated By'. The rule 'nw60045' is selected, showing a condition of 'service = 110 && tcp.dstport = I-109,111-u &&...'. The 'Alert' column for this rule shows 'alert.id'.

Status	Name	Condition	Session Data	Alert	Last Updated By
☐	nw110025	ip.proto=1 && (filename exists filetype exists)		alert.id	
☐	nw110030	service=0 && (filename exists filetype exists)		alert.id	
☐	nw60005	service = 53 && udp.dstport = I-52,54-u && stre...		alert.id	
☐	nw60010	service = 23 && tcp.dstport = I-22,24-u && strea...		alert.id	
☐	nw60015	service = 21 && tcp.dstport = I-20,22-u && strea...		alert.id	
☐	nw60020	service = 80 && tcp.dstport = I-79,81-u && strea...		alert.id	
☐	nw60025	service = 22 && tcp.dstport = I-21,23-u && strea...		alert.id	
☐	nw60030	service = 25 && tcp.dstport = I-24,26-u && strea...		alert.id	
☐	nw60035	service = 67 && tcp.dstport = I-66,68-u && strea...		alert.id	
☐	nw60040	service = 69 && tcp.dstport = I-68,70-u && strea...		alert.id	
☒	nw60045	service = 110 && tcp.dstport = I-109,111-u &&...		alert.id	
☐	nw60050	service = 119 && tcp.dstport = I-118,120-u && s...		alert.id	
☐	nw60075	service = 443 && tcp.dstport != 443,993,995 && ...		alert.id	
☐	nw60080	service = 520 && tcp.dstport = I-519,521-u && s...		alert.id	
☐	nw60085	service = 1433 && tcp.dstport = I-1432,1434-u &&...		alert.id	
☐	nw60090	service = 1521 && tcp.dstport = I-1520,1522-u &&...		alert.id	
☐	nw60095	service = 1719 && tcp.dstport = I-1718,1720-u &&...		alert.id	

At the bottom of the table, there are 'Apply' and 'Revert' buttons. The footer of the interface shows 'RSA NETWITNESS PLATFORM' and the version '11.5.0.0'.

4. Do one of the following:
 - If adding a new rule, click .
 - If editing a rule, select the rule from the rules list and click .
5. The Rule Editor Dialog is displayed with application rule parameters.



Rule Editor

Rule Definition

Rule Name

Condition

*All string literals and time stamps must be quoted.
Do not quote number values and ip addresses.
Examples : 1. device.group='Windows Compliance' && service = 443
2. time = '2015-jan-01 00:00:00' - u
3. ip.src = 10.0.0.0/8,172.16.0.0/12,192.168.0.0/16 || extension = 'torrent'*

Session Data

☒ Stop Rule Processing

☐ Keep

☐ Filter

☒ Truncate

☒ All

☐ After First Bytes

☐ After SSL/TLS Handshake

Session Options

☐ Alert ☐ Forward ☐ Transient

Alert On

NOTE: If applied to a session that is not SSL/TLS, this option will truncate the payload.

Reset Cancel OK

- a. In the **Rule Name** field, type a name for the rule. For example, for a rule that truncates all SMB, type **Truncate SMB**.
- b. In the **Condition** field, build the rule condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta from the window actions. As you build the rule definition, NetWitness Platform displays syntax errors and warnings. For example, to truncate all SMB, type **service=139**.

All string literals and time stamps must be quoted. Do not quote number values and IP addresses. [Configure Decoder Rules](#) provides additional details.

- c. If you want rule evaluation to end with this rule, check the **Stop Rule Processing** checkbox.
 - d. In the **Session Data** section, choose one of the following actions to apply when a matching packet is found:
 - **Keep**: The packet payload and associated meta are saved when they match the rule.
 - **Filter**: The packet is not saved when it matches the rule.
 - **Truncate**: Select a truncate option to execute when a packet matches the rule. The example uses the **All** option.
 - **Truncate All** to save the packet headers and associated metadata, and do not save the packet payload.
 - **Truncate After First <n> Bytes** to save the packet headers and associated metadata, and do not save the packet payload after the specified first <n> bytes, where <n> is a number of bytes.
 - **Truncate SSL/TLS Handshake** to truncates the payload for all sessions except in the case of an SSL/TLS session, where the SSL exchange is preserved, but the rest of the payload is not saved. This option is for use with SSL parsers.
 - e. In the **Session Options** section, do any of the following:
 - **To generate a custom alert** when a session metadata matches the rule, enable the Alert flag and select the name of the alert meta from the **Alert On** drop-down list.
 - **To perform syslog forwarding** when the log matches the rule, enable the **Forward** flag. Make sure that:
 - You have enabled both the Alert and Forward flags to carry out syslog forwarding.
 - The name of the rule mentioned in the Rule Editor dialog matches the syslog forwarding destination name specified in the Log Decoder > View > Explore > `/decoder/config/logs.forwarding.destination` parameter
 - **To prevent the alert metadata that is created from being written to the disk**, enable the **Transient** flag.
6. To save the rule and add it to the grid, click **OK**.

The rule is added at the end of the grid or inserted where you specified in the context menu. The plus sign is displayed in the **Pending** column.
 7. Check that the rule is in the correct execution sequence with other rules in the grid. If necessary, move the rule.
 8. To apply the updated rule set to the Decoder or Log Decoder, click **Apply**.

NetWitness Platform saves a snapshot of the currently applied rules, then applies the updated set to the Decoder and removes the pending indicator from the rules that were pending.

Monitor Application Rules

The Decoder and Log Decoder keep track of how many times each application rule matches a session. These stats can be viewed by connecting to the Decoder or Log Decoder Explore view and viewing the properties on the `/decoder/config/rules/application` folder. Then, send the command "statdump" to that folder. The output of this message is a listing of the number of times each application rule is hit. The listing is ordered in the same order as the contents of the rule definitions in the `/decoder/config/rules/application` folder. For example, on a system with three application rules:

```
0001: hits=6543 loaded=true
0002: hits=9294 loaded=true
0003: hits=43 loaded=true
```

The hit counters for the application rules are reset whenever the parsers are reloaded.

Configure Correlation Rules

Basic Correlation Rules are applied at the session level and alert the user to specific activities that may be occurring in their environment. NetWitness Platform applies correlation rules over a configurable sliding time window. When the conditions are met, alert metadata is created for this activity and there is a visible indicator of the suspicious activity.

The following are sample correlation rules illustrating two use cases and the syntax.

Objective: In sessions where `port.dst` exists, if there is any combination of `ip.src` and `ip.dst` where the count of unique instances of `port.dst` > 5 within one minute, then alert. To achieve this objective, create a rule as follows:

- Rule Name: IPv6 Vertical TCP Port Scan 5
- Rule: `port.dst exists`
- Instance Key: `ip.src, ip.dst`
- Threshold: `u_count(port.dst)>5`
- Time Window: 1 min

Objective: In sessions where `action==login` and `error==fail`, if there is any combination of `ip.src` and `ip.dst` that appears in more than 10 sessions within five minutes, then alert. To achieve this objective, create a rule as follows:

- Rule Name: IPv4 Potential Brute Force 10
- Rule: `action='login' && error='fail'`
- Instance Key: `ip.src, ip.dst`
- Threshold: `count()>10`
- Time window: 5 mins

Both sample rules have the same instance key: `ip.src` and `ip.dst`. Because we are looking for unique combinations of `ip.src` and `ip.dst` that match the correlation condition, **`ip.src` and `ip.dst` are primary keys.**

Threshold can include an **associated key** that identifies the meta type that we are counting to determine if the condition is satisfied. In the first example, the associated key specified in Threshold is `port.dst`. We are counting unique instances of `port.dst` for every `ip.src/ip.dst` pair. In the second example, the associated key is not specified in the Threshold because it is merely a count of sessions. It is helpful to think of this scenario as counting unique session IDs and the associated meta is implicitly `session.id`. We are counting unique `session.id` for every `ip.src/ip.dst` pair.

Invalid use case: In sessions where (rule), if there is any combination of `ip.src` and `ip.dst` that have a unique count of `ipv6.dst` > 5 within (time window), then alert. This case does not work because the associated key `ipv6.dst` is an IPv6 meta type. IPv4 and IPv6 meta types are not permitted to be used as associated keys.

To add or edit a correlation rule

1. Go to  (Admin) > **Services**, select a service, and click  > **View** > **Config**.
The Service Config view for the selected service is displayed.


2. Select the **Correlation Rules** tab.

The screenshot displays the RSA NetWitness Platform interface. The top navigation bar includes links for Investigate, Respond, Users, Hosts, Files, Dashboard, and Reports. Below this, a secondary navigation bar lists various categories: HOSTS, SERVICES, EVENT SOURCES, ENDPOINT SOURCES, HEALTH & WELLNESS, SYSTEM, and SECURITY. The SERVICES category is selected, and the 'Config' sub-menu is open, showing options like Change Service, Decoder, and Config. The 'Correlation Rules' tab is highlighted within the Config sub-menu. Below the tabs, there are controls for enabling/disabling rules and a filter box. A table with columns for Status, Pending, Name, Condition, Instance Key, Threshold, and Time Window is visible. At the bottom, there are 'Apply' and 'Revert' buttons. The footer shows 'RSA NETWITNESS PLATFORM' and the version '11.5.0.0'.

3. In the **Correlation Rules** tab, do one of the following:

- If adding a new rule, click **+**.

- If editing a rule, select the rule from the rules grid and click . The Rule Editor dialog is displayed with correlation rule parameters.



Rule Editor

Rule Definition

Rule Name

Condition

*All string literals and time stamps must be quoted.
Do not quote number values and ip addresses.
Examples : 1. device.group='Windows Compliance' && service = 443
2. time = '2015-jan-01 00:00:00' - u
3. ip.src = 10.0.0.0/8,172.16.0.0/12,192.168.0.0/16 || extension = 'torrent'*

Correlation Fields

Threshold

Instance Key

Time Window

Reset Cancel OK

4. In the **Rule Name** field, type a name for the rule. For example, to create the sample rule, **IPv6 Vertical TCP Port Scan 5**.
5. In the **Condition** field, build the rule condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta from the window actions. As you build the rule definition, syntax errors and warnings are displayed by NetWitness Platform. For example, to create the sample rule, type **tcp.dstport exists**. When this condition is matched, the session data action is performed.
All string literals and time stamps must be quoted. Do not quote number values and IP addresses. [Configure Decoder Rules](#) provides additional details.
6. In the **Threshold** field, use one of the threshold parameters to specify the minimum number of occurrences required to create a correlation session and an associated key if required. The associated key cannot be an IPv4 or IPv6 meta type.
 - `u_count(associated_key)` = the count of unique values of the specified key
 - `sum(associated_key)` = the values of the specified key
 - `count` = number of sessions (no associated key is specified)
7. In the **Instance Key** field, select the target indicator to base the event upon. This can be a single key or a compound key (two primary keys, separated by a comma).

8. In the **Time Window**, set the duration during which the threshold must be reached to create a correlation session.
9. To save the rule and add it to the grid, click **OK**.
The rule is added at the end of the grid or inserted where you specified in the context menu. The plus sign is displayed in the **Pending** column.
10. Check that the rule is in the correct execution sequence with other rules in the grid. If necessary, move the rule.
11. To apply the updated rule set to the service, click **Apply**.
NetWitness Platform saves a snapshot of the currently applied rules, then applies the updated set to the Decoder or Log Decoder.

Configure Network Rules

Network rules are applied at the packet level on a Decoder and are made up of rule sets from Layer 2, Layer 3, and Layer 4. Multiple rules can be applied at the packet level to a Decoder. Network rules can apply to multiple network layers (for example, when a network rule filters out specific ports for a specific IP address). Network rules do not apply to Log Decoders, they apply only to Network Decoders.

You can create and manage network rules in the Services Config view > Network Rules tab.

Note: Because network rules are applied on the packet level, you must specify both the source and destination meta keys in the rule condition. This is because the packet flow can occur on both sides at the packet level, while the directions are still not determined. For example, if you want to filter traffic on port 553 and port 5553, the condition should be written as follows:

```
port.src=553 || port.dst=553 || port.src=5553 || port.dst=5553
```

You must also specify both the source and destination meta keys in the conditions for `ip.src` and `ip.dst`. Specifying only source or destination in the condition will not work as expected.

Supported Meta Keys in Network Rule Conditions

The following table describes the meta keys that NetWitness Platform supports for use in network rule conditions.

Meta Key	Description
<code>eth.addr</code>	Ethernet source or destination address. Commonly known as the MAC address.
<code>eth.dst</code>	Destination Ethernet address. This is the same as the Ethernet address field except that it selects only packets where the destination address matches the selected value (s).
<code>eth.src</code>	Same as Ethernet destination except that it focuses on the source address.
<code>eth.type</code>	Ethernet frame type.
<code>hdlc.type</code>	Frame type of the HDLC frame.
<code>alias.ip</code>	IPv4 source or destination address in standard form. IP addresses can be entered in CIDR notation for subnets.
<code>ip.dst</code>	Destination IPv4 address in standard form. IP addresses can be entered in CIDR notation for subnets.
<code>ip.proto</code>	IPv4 protocol field.
<code>ip.src</code>	Source IPv4 address in standard form. IP addresses can be entered in CIDR notation for subnets.
<code>alias.ipv6</code>	IPv6 source or destination address in hex format. Generally IPv6 addresses are written as eight groups of four hex digits, thus expressing the entire 128 bit address length. Supports notation to represent multiple blocks of 0000 in an address. Does not support CIDR notation.

Meta Key	Description
<code>ipv6.dst</code>	Destination IPv6 address in hex format.
<code>ip.proto</code>	IPv6 protocol field. This maps to the Next Header field in the IPv6 header and uses the same values as the IPv4 protocol field.
<code>ipv6.src</code>	Source IPv6 address in hex format.
<code>port.dst</code>	Destination TCP port.
<code>tcp.port</code>	TCP source or destination port.
<code>port.src</code>	Source TCP port.
<code>port.dst</code>	Destination UDP port.
<code>udp.port</code>	UDP source or destination port.
<code>port.src</code>	Source UDP port.

The following are sample network rules.

To truncate all SSL from the source port, create a rule as follows:

- Rule Name: Truncate SSL
- Condition: `port.src=443`
- Rule Action: Truncate

To filter subnet traffic, create a rule as follows:

- Rule Name: Subnet Filter
- Condition: `alias.ip=192.168.2.0/24`
- Rule Action: Filter

Meta entities, which provide a way to work with several meta keys at the same time, can be used in application rules, but are not supported in network rules as the metadata available are too limited. For more information on meta entities, see the *Core Database Tuning Guide*.

To add or edit a network rule:

1. Go to  (Admin) > **Services**, select a Decoder service, and  > **View** > **Config**. The Services Config view for the selected service is displayed.

2. Select the **Network Rules** tab.

The Network Rules tab is displayed.


The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. Below this, a secondary bar shows 'HOSTS', 'SERVICES', 'EVENT SOURCES', 'ENDPOINT SOURCES', 'HEALTH & WELLNESS', 'SYSTEM', and 'SECURITY'. The 'SERVICES' section is expanded, showing 'NH - Decoder' and 'Config'. The 'Config' tab is active, displaying a list of network rules. The rules are listed in a table with columns for Status, Name, Condition, Packet Data, and Alert. Two rules are visible: 'upgrade1' with condition 'tcp.srcport=443' and alert 'keep', and 'upgrade2' with condition 'ip.src=127.0.0.1' and alert 'truncate'. At the bottom of the table, there are 'Apply' and 'Revert' buttons. The footer of the interface shows 'RSA NETWITNESS PLATFORM' and the version '11.5.0.0'.

Status	Name	Condition	Packet Data	Alert
☒	upgrade1	tcp.srcport=443	keep	
☒	upgrade2	ip.src=127.0.0.1	truncate	

3. In the **Network Rules** tab, do one of the following:

- If adding a new rule, click **+**.

- If editing a rule, select the rule from the rules list and click . The Rule Editor dialog is displayed.



4. In the **Rule Name** field, provide a name for the rule. For example, for a rule that truncates all SSL from the source port, type **SSL Truncate**.
5. In the **Condition** field, build the rule condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta values from the window actions. As you build the rule definition, NetWitness Platform displays syntax errors and warnings. For example, to truncate all SSL from the source port, **tcp.srcport=443**.
All string literals and time stamps must be quoted. Do not quote number values and IP addresses. [Configure Decoder Rules](#) provides additional details. [Supported Meta Keys in Network Rule Conditions](#) describes the meta keys that NetWitness Platform supports for use in network rule conditions.
6. If you want rule evaluation to end with this rule, select the **Stop Rule Processing** checkbox.
7. In the **Session Data** section, choose one of the following actions to apply when a matching packet is found:
 - **Keep**: The packet payload and associated metadata are saved when they match the rule.
 - **Filter**: The packet is not saved when it matches the rule.
 - **Truncate**: The packet payload is not saved when it matches the rule, but packet headers and associated metadata are retained.
8. In the **Session Options** section, select all options that apply of these four.
 - **Assemble**: The assembler assembles the packet chain when it matches the rule.
 - **Network Meta**: The packet generates network metadata when it matches the rule.
 - **Application Meta**: The packet generates application metadata when it matches the rule.
 - **Alert**: The packet generates a custom alert when metadata matches the rule.

9. To save the rule and add it to the rules list, click **OK**.
The rule is added at the end of the list or inserted where you specified in the context menu.
10. Check that the rule is in the correct execution sequence with other rules in the list. If necessary, move the rule.
11. To apply the updated rule set to the Decoder, click **Apply**.
NetWitness Platform saves a snapshot of the currently applied rules, then applies the updated set to the Decoder and removes the pending indicator from the rules that were pending.


Fix Rules with Invalid Syntax

After an update to NetWitness Platform 11.x, the user interface highlights any rules with invalid syntax. The Rule Editor provides additional tooltips. After you fix the rules, the highlights disappear. [Configure Decoder Rules](#) provides guidelines that all queries and rule conditions in NetWitness Platform must follow.

To correct rules with invalid syntax:

1. Go to  (Admin) > **Services**.
2. In the **Services** view, select a Decoder and  > **View** > **Config**.
3. In the **Services Config** view, select one of the Rules tabs: Network Rules, App Rules, or Correlation Rules.

The Rules tab for the selected rule type shows the number of rules using invalid syntax and the invalid rules are highlighted.




The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes tabs for Investigate, Respond, Users, Hosts, Files, Dashboard, and Reports. Below this, the 'SERVICES' tab is selected, and the 'NH - Decoder' is chosen. The 'Config' view is active, showing a list of rules under the 'App Rules' tab. A yellow banner at the top of the rule list states '1 rule is using deprecated syntax.' The table below lists rules with columns for Status, Name, Condition, Session Data, Alert, and Last Updated By. The 'DecTester' rule is highlighted in yellow, indicating it has deprecated syntax. The condition for 'DecTester' is 'ip.src = "10.30.30.30"'. Other rules include 'nw110025', 'nw110030', 'nw60005', 'nw60010', 'nw60015', 'nw60020', 'nw60025', 'nw60030', 'nw60035', 'nw60045', 'nw60050', 'nw60055', 'nw60060', 'nw60065', and 'nw60095'.

Status	Name	Condition	Session Data	Alert	Last Updated By
☐	nw110025	ip.proto=1 && (filename exists filetype exists)		alert.id	
☐	nw110030	service=0 && (filename exists filetype exists)		alert.id	
☐	nw60005	service = 53 && udp.dstport = I-52,54-u && stre...		alert.id	
☐	nw60010	service = 23 && tcp.dstport = I-22,24-u && strea...		alert.id	
☐	nw60015	service = 21 && tcp.dstport = I-20,22-u && strea...		alert.id	
☐	nw60020	service = 80 && tcp.dstport = I-79,81-u && strea...		alert.id	
☐	nw60025	service = 22 && tcp.dstport = I-21,23-u && strea...		alert.id	
☐	nw60030	service = 25 && tcp.dstport = I-24,26-u && strea...		alert.id	
☐	nw60035	service = 67 && tcp.dstport = I-66,68-u && strea...		alert.id	
☐	DecTester	ip.src = "10.30.30.30"		alert.id	
☐	nw60045	service = 110 && tcp.dstport = I-109,111-u && s...		alert.id	
☐	nw60050	service = 119 && tcp.dstport = I-118,120-u && s...		alert.id	
☐	nw60055	service = 135 && tcp.dstport = I-134,136-u && s...		alert.id	
☐	nw60060	service = 137 && tcp.dstport = I-136,138-u && s...		alert.id	
☐	nw60065	service = 139 && tcp.dstport = I-138,140-444,44...		alert.id	
☐	nw60095	service = 1719 && tcp.dstport = I-1718,1720-u &&...		alert.id	

At the bottom of the table, there are 'Apply' and 'Revert' buttons. The footer of the interface shows 'RSA NETWITNESS PLATFORM' and the version '11.5.0.0'.

4. Select an invalid rule and click .

The Rules Editor shows additional information for the invalid rule.



The screenshot shows the 'Rule Editor' window. The 'Rule Definition' section has 'Rule Name' set to 'test' and 'Condition' set to 'ip.src = "10.0.0.0/8"'. A red box highlights the condition field, and a red error icon is next to it. A tooltip message reads: 'rule 'temp': 9: expected a comma-separated list of IPv4 address ranges, values, or a CIDR range or a comma-separated list of keys'. Below the condition field, there is a note: 'All string literals and time stamps must be quoted. Do not quote number values and ip addresses. Examples : 1. ip.src = 10.0.0.0/8,172.16.0.0/12,192.168.0.0/16 2. tcp.srcport= 20,21,22,80'. The 'Session Data' section has 'Stop Rule Processing' checked, 'Keep' selected, and 'Filter' and 'Truncate' unselected. The 'Session Options' section has 'Assemble', 'Application Meta', and 'Network Meta' checked, and 'Alert' unselected. At the bottom are 'Reset', 'Cancel', and 'OK' buttons.

5. In the **Condition** field, correct the rule syntax.

All string literals and time stamps must be quoted. Do not quote number values and IP addresses. [Configure Decoder Rules](#) provides additional details.

For example, if the invalid rule condition is `ip.src="10.30.30.30"`, correct the syntax by removing the quotes: `ip.src=10.30.30.30`

6. Do one of the following:
- To correct the rule individually, click **Save**.
The corrected rule is applied independently to the Decoder. The corrected rule appears on the Rules tab without highlights.
 - To correct the rule and apply the rule to the Decoder later with other rules, click **OK**.
The corrected rule appears on the Rules tab without highlights. The rule is not applied to the Decoder.

Decoder Commands for Managing Rules

In the NetWitness Core database, the Rules tree holds the main functionality related to managing rules for all Core services that have rules: Concentrators, Decoders, Log Decoders, and Archivers. Although you can manage rules in the NetWitness Platform user interface, advanced users may prefer to manage rules using a command line to add, merge, replace, delete, and validate rules on a service. This section provides a brief overview of the commands and their usage. These are the available commands:

- `add` - Adds a single rule at the specified position.
- `clear` - Deletes all existing rules in the current node on the service. For example, using the command in `/decoder/config/rules/application` node deletes all existing application rules on the Decoder.
- `delete` - Deletes one or more rules at a specified position and count.
- `merge` - Merges a pushed rule set with an existing rule set. Existing rules that match the incoming rules (by name or rule) are replaced; otherwise, rules are inserted by the position indicated as described in [merge Command](#).
- `replace` - Deletes all existing rules and replaces them with the incoming rule set.
- `validate` - Validates the syntax of a rule, but does not validate the meta keys.

add Command

The `add` command adds the rule to the existing rule set. Formatting is important because the API uses double quotes in the rule language and also uses double quotes as parameters to all RSA NetWitness Platform APIs. Therefore, you must escape any double quotes in the rule itself by preceding it with a backslash (`\`) character. This is the syntax of the command:

```
add rule=<string> name=<string> alert=<string, optional> atPos=<uint32, optional>
```

- `rule` is the rule to add. Be sure to place double quotes around any rules with white space and to escape any double quotes that are part of the rule with a backslash.
- `name` is the name of the rule.
- `alert` is the alert for the rule (if any).
- `atPos` is the position at which the rule should be added (1 based). Zero is the top of the list and any number larger than the current size of the list is appended to the list.

This is an example of command to add a rule using `NwConsole`

```
send /decoder/config/rules/application add rule="ip.src exists" order=1  
alert=ioc name=testrule
```

For example, take the following rule:

```
alias.host = "myPC" && country.src="china","russian federation"
```

To add this as a rule, you would need to send the parameters as follows:

```
rule="alias.host = \"myPC\" && country.src=\"china\", \"russian federation\""  
name=myRule filter
```

Notice how all the double quotes had to be escaped inside the rule parameter. A simple trick to make this more readable is to use single quotes inside the rule. Single and double quotes are interchangeable in the rule and query language, but not in parameters for the API (only double quotes are supported there). Therefore, this is more readable:

```
rule="alias.host = 'myPC' && country.src='china','russian federation'"
name=myRule filter
```

merge Command

The `merge` command is used to merge an incoming list of rules with the existing rules on the service. This is how it works:

- It finds existing rules that match via the name OR via a matching rule, updates the existing rule name, and keeps the same position.
- It inserts new rules into the rule list based on the NUMBER position. If the number is zero, it goes to the top of the list.
- It processes the rules in the order received so if you have two rules numbered zero, the second rule is processed after the first and claims the top spot. All existing rules are pushed down two places. Any numbers higher than existing rule positions are appended after the last existing rule and numbered in sequence.
- Any non-numbered rule is appended after the last existing rule and numbered in sequence.

This is the syntax of the merge command:

```
merge --file-data=<string> --file-format<string>
```

- `file-data` is the full path and name of the rules file to merge.
- `file-format` is the format of the rules file. Valid values are `params-list`, `string`, `params`, `binary`, and `params-binary`.

Methods of Sending a List of Rules to a Service

There are two ways to send a list of rules. You can send them as a `.nwr` (NetWitness Rule) file or as a numbered set of parameters, each number indicates the position to insert the rule at as well as the encoded rule. If you want to see the current list of rules on a service, you need to run the `ls` command on the rule category (for instance, application rules on a Decoder are found in `/decoder/config/rules/application`).

This is an example of commands to list the existing rules using NwConsole:

```
login <hostname>:50004 <username> <password>
cd /decoder/config/rules/application
ls
```

This is another example to list existing rules in NwConsole:

```
send /decoder/config/rules/application ls
```

This is an example of the command to point to network rules in the RESTful port, which supports a basic admin HTML app.

```
http[s]://<decoder>:50104/decoder/config/rules/network
```

Send a NetWitness Rule File

Let's start with an example `nwr` file, each rule must be on a separate line:

```
rule="ip.src=192.168.0.1" name=first keep
rule="ip.src=192.168.1.1" name=second alert=ioc
rule="ip.src=192.168.2.1" name=third filter
```

To push and merge rules using NwConsole, use the following commands:

```
login <hostname>:50004 <username> <password>
send /decoder/config/rules/application merge --file-data=/root/App_
Rules.nwr --file-format=params-list
```

To replace the existing rules with the rules in the file, instead of using the `merge` command, use the `replace` command.

```
send /decoder/config/rules/application replace --file-
data=<pathname> --file-format=params-list
```

To merge the rules in an `nwr` file using the RESTful port, you can use a `curl` command that pushes the rules:

```
curl -u "<username>:<password>" -H "Content-Type: application/octet-
stream" --data-binary @<pathname> -X POST
"http://<hostname>:50104/decoder/config/rules/application?msg=merge"
```

The examples are pushing application rules. To push network rules, send the rules to `/decoder/config/rules/network`. For correlation rules, send the rules to `/decoder/config/rules/correlation`.

Send Numbered Parameters

The other way to send a list of rules is to send them as numbered parameters. The difficulty with this method is remembering to escape the quotes within each numbered rule. Though it is only a problem if you are trying to do it by hand. For instance, to send the same rules above as parameters via NwConsole, use the following command:

```
send /decoder/config/rules/application merge
1="rule=\"ip.src=192.168.0.1\" name=first keep"
2="rule=\"ip.src=192.168.1.1\" name=second alert=ioc"
3="rule=\"ip.src=192.168.2.1\" name=third filter"
```

This command is hard to read because you have to escape the inner quotes with a backslash (`\`). Otherwise, these two commands accomplish the same thing. Merging or adding three rules in positions 1, 2 and 3. If you think the above was hard to read, this is what the equivalent `curl` command looks like:

```
curl -u "<username>:<password>"
"http://<hostname>:50104/decoder/config/rules/application?msg=merge&1=rule%3D%
22ip.src%3D192.168.0.1%22%20name%3Dfirst%20keep&2=rule%3D%22ip.src%3D192.168.1
.1%22%20name%3Dsecond%20alert%3Dioc&3=rule%3D%22ip.src%3D192.168.2.1%22%20name
%3Dthird%20filter"
```

For more details on how to escape double quotes inside parameters, see [add Command](#).

Ordering Rules When Pushing

Pushed rules are ordered in one of two ways. When passing as parameters, the number of each parameter determines the insertion order. If it is not actually a number, merge checks for an `order` parameter within the rule itself and uses that value if found.

Note: Using `order` is the only way to set the order with a `.nwr` file. If neither a number nor an `order` parameter is found, there are no guarantees of the insertion order.

Example

A Decoder has the following application rules installed; notice the numbering is ALWAYS consecutive and starts at 1:

```
0001 : rule="ip.src = 192.168.0.1 || ip.dst = 192.168.0.1 || alias.host = 'My-PC'" name=first keep
0002 : rule="ip.src=192.168.1.1" name=second alert=ioc
0003 : rule="ip.src=192.168.2.1" name=third filter
```

And you want to merge the following four rules:

```
rule="ip.src=192.168.3.1" name=third keep
rule="ip.dst=192.168.4.1" name=NewRule filter order=0
rule="alias.host = 'pc1','pc2'" name=filterTheseNames filter order=append
rule="service=80,443" name=web filter order=3
```

Use any method to push your rules and this is what you end up with:

```
0001 : rule="ip.dst=192.168.4.1" name=NewRule filter order=1
0002 : rule="ip.src = 192.168.0.1 || ip.dst = 192.168.0.1 || alias.host = 'My-PC'" name=first keep order=2
0003 : rule="service=80,443" name=web filter order=3
0004 : rule="ip.src=192.168.1.1" name=second alert=ioc order=4
0005 : rule="ip.src=192.168.3.1" name=third keep order=5
0006 : rule="alias.host = 'pc1','pc2'" name=filterTheseNames filter order=6
```

Are there any surprises here? This is how each rule was processed.

1. rule="ip.src=192.168.3.1" name=third keep

This rule had the same name as an existing rule on the Decoder (third). So the rule updated the existing rule, changing `_filter_` to `_keep_`.

2. rule="ip.dst=192.168.4.1" name=NewRule filter order=0

This rule is new and had `order=0` in it, which means insert at the very top.

3. rule="alias.host = 'pc1','pc2'" name=filterTheseNames filter order=append

This rule had a non-number `append` for `order`, therefore, it went to the end of the list. You can accomplish the same thing by giving a very large number, like `999999`.

4. rule="service=80,443" name=web filter order=3

This rule is last but has `order=3`, therefore, if it does not match an existing rule by name or the text of the rule itself, it should be placed in position 3. And there it is, the third rule in the list. Any rules that follow were pushed further down.

replace Command

The `replace` command removes all existing rules and replaces them with the incoming rule list. Refer to [merge Command](#) for details on how to format the incoming rule list and how ordering works.

This is an example of the `replace` command using a NetWitness Rule File :

```
send /decoder/config/rules/application replace --file-data=/root/Decoder-AppRules.nwr --file-format=string
```

This is an example of the `replace` command using Numbered Parameters :

```
send /decoder/config/rules/application replace 1="rule=\"ip.src exists\" name=\"test rule\" order=1 alert=ioc"
```

clear Command

The `clear` command removes all existing rules on the service. This is an example of the command:

```
send /decoder/config/rules/application clear
```

delete Command

The `delete` command deletes one or more rules on the service.

```
delete atPos <uint32> count <uint32, optional>
```

- `atPos` deletes the rule at the given position. Rules are numbered starting with 1 and go in sequential order.
- `count` deletes one or more rules starting `atPos`. This is an optional parameter defining the number of rules to delete starting `atPos`. The default value is 1.

This example of the command deletes four rules beginning at position 0003:

```
send /decoder/config/rules/application delete atPos=0003 count=4
```

validate Command

The `validate` command takes the provided rule and verifies that it parses correctly. Keep in mind that this command cannot verify whether language keys and entities are valid.

```
validate rule <string>
```

`rule` - is the name of the rule to validate. Make sure to place double quotes around any rules with white space.

Configure Parsers and Feeds

Parsers and feeds are responsible for analyzing the packets and logs when captured or imported in Decoders. Most commonly, they are used for static metadata extraction and service identification. The flexible definition allows custom extension of the core defined services to provide extra service type identification and metadata extraction. This is important due to the volume of custom applications that are used on networks.

See the following sections for details about configuring parsers and feeds.

- [Configure Parsers](#)
- [Configure Feeds](#)

Configure Parsers

NetWitness Platform has a set of native parsers that are defined by the system, and also provides the option to add additional parsers. Each parser is configurable in the [Services Config View - General Tab](#). The Parser Configuration panel provides a way to enable or disable parsers to use on Decoders in addition to limiting the metadata that the parser creates.

There are also several types of custom configurable parsers:

- GeoIP2 – This parser associates IP addresses with geographical locations. For new installations and upgrades, the GeoIP2 parser is enabled by default. For more information on these parsers, see [GeoIP2 Parsers](#).
- Search – This parser is user-configured to generate metadata by scanning for pre-defined keywords and regular expressions.
- FLEXPARSE (deprecated) – This is a generic parser definition language for extending the existing application protocol support of the Decoder. By default this parser is disabled (see [Enable or Disable Lua and Flex Parsing Systems](#)).
- Lua – This parser is defined using the Lua scripting language for extending the existing application protocol support of the Decoder.
- Log – This application parser supports the Log Decoder and is configured to generate metadata by scanning log files.
- Snort® – This parser supports the payload detection capabilities of Snort IDS rules. Snort rules and configuration are added to the `parsers/snort` directory for Investigation and Decoder (see [Decoder Snort Detection](#)).

In the Services Config view > Parsers tab, you can view deployed parsers on a Decoder, upload parsers, and delete deployed parsers. The user interface includes an Indicator if the parser originated from Live Services, installed through NetWitness Platform, or uploaded manually. Parsers can be added and removed while a Decoder is running without affecting capture.

Note: To pass options to parsers, you must first give the name of the parser and then the options to be passed in this format:

```
<ParserName>=<ParserOptions>"<Whitespace><ParserName2>=<Parser2Options>"
```

Each ParserName=Value option must be separated by whitespace. Normally, the Value must have double quotes around it. The Value itself can sometimes list multiple Option=Value pairs, each separated by whitespace, and if those values have whitespace, they must be in escaped double quotes. To escape a quote, place a backslash before it: \".

This is an example of defining options for Parser1, Parser2, and Parser3:

```
Parser1="Option1=\"Option1 Value With Space\" Option2=Option2ValueNoSpace"
Parser2="Option1=Value" Parser3="op1=val1 op2=val2 op3=\"another value\""
```

In addition, you can download parsers using NetWitness Platform Live Services.

Upload and Delete Custom Parsers

RSA NetWitness Platform has the ability to upload parsers from your local system and delete these parsers.


Upload Parsers to a Decoder or Log Decoder

The Upload option in the Service Config view > Parsers tab displays the Upload Parsers dialog, in which you can manage the uploading of parsers to a Decoder or Log Decoder. In the File list, you prepare a list of parsers for uploading. You can add files from a directory structure, and delete files from the list if you decide that you don't want to upload a particular file. When the list is ready, clicking Upload starts the upload process.

1. Go to  (Admin) > **Services**, select a service, and click  > **View > Config**.
The Config view for the selected service is displayed.
2. Click the **Parsers** tab.

- Click  **Upload**.

The Upload Parsers dialog is displayed.




- Click .

A file selection dialog is displayed.


- Select the **.flex**, **.parser**, and **.lua** files to be updated, and click **Open**.

The dialog closes, and the selected files are displayed in the File list.



6. Click **Upload**.

The Upload Job grid shows the progress of the upload jobs with each job representing a file being uploaded.



7. Use any of the Upload grid tools to manage the upload of selected jobs: pause and resume, cancel, and delete.

Once a job is complete, it is deployed on the Decoder and listed with the deployed parsers in Parsers tab.

Manage Upload Jobs

You can use any of the Upload grid tools to manage the upload of selected jobs: pause, resume, cancel, and delete.

- To cancel uploading a set of parsers while the upload is in queue or progress, click **Cancel**.
- To pause uploading a set of parsers, if the upload is not yet complete, click **Pause**.
- To resume uploading a set of parsers after a pause, click **Resume**.
- To delete an upload job, click .


Delete Deployed Parsers

The Delete option in the Service Config view > Parsers tab provides a way to delete deployed parsers from a Decoder or Log Decoder. Parsers can be added and removed while a Decoder is running without affecting capture.

Note: Unless otherwise stated, any reference to Decoders applies to Log Decoders as well.

To delete a parser from a Decoder:

1. Go to  (Admin) > **Services**, select a Decoder, and  > **View** > **Config**.
The Services Config view for the selected service is displayed.
2. Click the **Parsers** tab.



3. In the **Parsers** tab, select one or more parsers to delete.
4. Click .
A dialog requests confirmation that you want to delete the parsers.
5. If you want to delete the parsers, click **Yes**.
The parsers are removed from the Decoder immediately.

Enable and Configure the Entropy Parser

Beginning with NetWitness Platform 11.0, the administrator can configure a Decoder to use a NetWitness native parser, known as the Entropy parser. When the Entropy parser is enabled, analysts have visibility into channels that are trying to blend in with other traffic, but do not follow normal protocol behavior. This helps to identify channels that do not conform to the normal environment traffic baseline, and may be worthy of investigation.

The parser creates meta keys, based on statistics collected by the native NetWitness Platform parser, that help to identify behavior of any channel that is getting lots of network traffic. When the parser is first enabled, the analyst needs to become familiar with overall behavior for the different channels seen in a captured session to understand the frequency of bytes and the normal client and server payload. Once the normal behavior is known, analysts can use the meta keys to find behavior that does not match the expected.

By default, the Entropy parser generates 10 additional meta keys that do not add significantly to the load on a Decoder, and are useful for this specialized case. The parser is disabled by default.

Enable indexing if you have interest in exploring interesting sessions based on payload byte analysis of the packets. By default, to make indexing easier, the normal `Float32` value for `entropy.req` and `entropy.res` is multiplied by 10k and stored in a `UInt16` (thus giving four digits of precision, 0 to 10,000).

However, if you define the `entropy.*` fields in the Decoder language to be `Float32`, the Decoder will store it as a float with a range of 0.0 to 1.0. Take care to change the language everywhere if you decide to keep it as a `Float32`.

RSA does not recommend indexing as a `Float32` because of the high unique counts due to minute changes in precision.

These are the new meta keys generated by the Entropy parser by default:

- `entropy.req` and `entropy.res`: These meta keys capture entropy using the Shannon entropy equation, which has a floating point value as a result. The floating point value of 0 to 1.000 is multiplied by 10000 and written in NetWitness Platform as `UInt 16`, an unsigned integer of 0 through 10000. .
- `mcb.req` and `mcb.res`: The most common byte is simply which byte for each side (0 thru 255) was seen the most.
- `mcbc.req` and `mcbc.res`: The most common byte count is the number of times the most common byte (above) was seen in the session streams.
- `ubc.req` and `ubc.res`: - Unique byte count is the number of unique bytes seen in each stream. 256 would mean all byte values of 0 thru 255 were seen at least once.

To enable and configure the Entropy parser on a Decoder:

1. Go to  (Admin) > **Services**, select a Decoder service and click   **View > Config**.
The Services Config view for the selected Decoder is displayed.
2. The Entropy parser is disabled by default. Click the drop-down list for **Entropy** in the **Config Value** column, and select **Enabled**. If you want to disable some of the meta keys, click the drop-down list

and select **Disabled** next to the meta key.


The screenshot shows the RSA NetWitness Platform configuration interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. Below this, a secondary navigation bar lists 'HOSTS', 'SERVICES', 'EVENT SOURCES', 'ENDPOINT SOURCES', 'HEALTH & WELLNESS', 'SYSTEM', and 'SECURITY'. The main content area is titled 'Config' and contains three tabs: 'General', 'Files', and 'Data Retention Scheduler'. The 'General' tab is active, showing 'System Configuration' and 'Decoder Configuration' sections. The 'Parsers Configuration' section is also visible, showing a list of parsers and their status. The 'Entropy' parser is highlighted with a red box and its status is 'Disabled'.

Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☒
SSL Port	56004
Stat Update Interval	1000
Threads	20

Name	Config Value
Adapter	Berkeley Packet Filter
Capture Interface Selected	packet_mmap_eth0 (bpf)
Cache	Cache Directory: /var/netwitness/decoder/cache
Cache Size	4 GB
Capture Settings	Assembler Maximum Size: 32 MB
Assembler Minimum Size	0

Name	Config Value
ALERTS	Enabled
DHCP	Enabled
DNS	Enabled
Entropy	Disabled
entropy.req	Disabled
entropy.res	Disabled
mcb.req	Disabled
mcb.res	Disabled
mcbc.req	Disabled
mcbc.res	Disabled
payload.req	Disabled
payload.res	Disabled
ubc.req	Disabled
ubc.res	Disabled
FeedParser	Enabled
FTP	Enabled

- Click **Apply**.
The Entropy parser is enabled and begins creating the new meta keys as configured in the Concentrator custom index file.
- In the Service Config view select the Concentrator that is aggregating traffic from this Decoder. Select **View > Files** and open the Custom Index file for the Concentrator. Look for the Entropy parser meta keys to see if they are included and uncommented.
By default the keys are commented out and therefore not enabled. To enable that part of the language the administrator needs to copy that part of index file into the `index-concentrator-custom.xml` and uncomment the key description line for each meta key. An example of the custom index file with the Entropy parser keys and instructions is shown below in [Entropy Parser Configuration in the Concentrator Custom Index File](#).
- With the Entropy meta keys enabled, they are available to analysts in Investigate, but hidden by default. To make the meta keys visible in the Investigate Values view, edit the default meta keys in the Default Meta Keys dialog so that they are open instead of hidden. You can manage these meta key the same way you manage other meta keys.



Entropy Parser Configuration in the Concentrator Custom Index File

The following is an excerpt of the Concentrator Index file lines that the administrator must copy to the custom index file. The comments provide guidance on configuring the parser.

```
<!-- This section is commented out because it's only used by the Entropy
parser which is disabled by default. To enable this part of the language, copy
to index-concentrator-custom.xml and uncomment the keys. HOWEVER, take note
that depending on how the Entropy parser is configured, the entropy.req and
entropy.res format might be a Float32 instead of a UInt16. So make sure to
change to the correct type if necessary.-->
```

```
<!-- Entropy parser meta - enable indexing if you have interest in exploring
this for interesting sessions based on payload byte analysis of the packets.
By default, to make indexing easier, the normal Float32 value for entropy.req
and entropy.res is multiplied by 10k and stored in a UInt16 (thus giving 4
digits of precision, 0 to 10,000). However, if you define the entropy.* fields
in the Decoder language to be Float32, it will store it as a float with a
range of 0.0 to 1.0. Take care to change the language everywhere if you decide
to keep it as a Float32. We do not recommend indexing as a Float32 because of
the high unique counts due to minute changes in precision. -->
```

```
<!--
```

```
<key description="Entropy Request (Client)" format="UInt16" level="IndexNone"
name="entropy.req" valueMax="10001"/>
```

```
<key description="Entropy Response (Server)" format="UInt16" level="IndexNone"
name="entropy.res" valueMax="10001"/>
```

```
-->
```

```
<!-- The most common byte is simply which byte for each side (0 thru 255) was
seen the most -->
<!--
<key description="Most Common Byte Request" format="UInt8" level="IndexNone"
name="mcb.req"/>
<key description="Most Common Byte Response" format="UInt8" level="IndexNone"
name="mcb.res"/>
-->

<!-- The most common byte count is the number of times the most common byte
(above) was seen in the session streams -->
<!--
<key description="Most Common Byte Count Request" format="UInt32"
level="IndexNone" name="mcbc.req" valueMax="500000"/>
<key description="Most Common Byte Count Response" format="UInt32"
level="IndexNone" name="mcbc.res" valueMax="500000"/>
-->

<!-- Unique byte count is the number of unique bytes seen in each stream. 256
would mean all byte values of 0 thru 255 were seen at least once -->
<!--
<key description="Unique Byte Count Request" format="UInt16" level="IndexNone"
name="ubc.req"/>
<key description="Unique Byte Count Response" format="UInt16"
level="IndexNone" name="ubc.res"/>
-->

<!-- The payload size metrics are the payload sizes of each session side at
the time of parsing. However, in order to keep indexing from having high
unique counts (bad for performance), the two payload size metas below are
indexed in buckets. -->
<!--
<key description="Payload Size Request" format="UInt32" level="IndexNone"
bucket="true" name="payload.req" valueMax="500000"/>
<key description="Payload Size Response" format="UInt32" level="IndexNone"
bucket="true" name="payload.res" valueMax="500000"/>
-->
```

Flex Parsers

There are two kinds of Flex parsers:

- **Service identification based solely on port.** These are parsers that use only the source or destination ports to identify the session application type (service). These are the most basic and easiest to define.
- **Service identification based on a found token(s).** These parsers use tokens to identify the service type. This is also an easy way to expand which service types are identified. These are important when identifying non-internet standard applications. These parsers require that the protocol has a definable token that can uniquely identify the service type.

Five common parser operations are:

- Match Port and Identify Immediately
- Match Port and Delay Identification
- Match Token and Identify Immediately
- Match Multiple Tokens
- Match Token and Create Metadata

Detailed language information and samples are provided in this topic. This topic describes the XML schema used to define a FlexParse file. The SML node, attribute, and values referenced in descriptive text are **bold**. The root node of every file must be the **parsers** node. Under that node there can be any number of parser nodes. Each parser node defines a single parser. A parser node can have an optional **declaration** node and any number of **match** nodes.

Topics

- [Arithmetic Functions](#)
- [Common Parser Operations](#)
- [General Functions](#)
- [Logging Functions](#)
- [Nodes](#)
- [Payload Functions](#)
- [Regex](#)
- [String Functions](#)

Arithmetic Functions

This topic defines language for the flex parser arithmetic functions.

This topic defines language for the flex parser arithmetic functions. All numbers are 64-bit unsigned values and subject to both underflow and overflow, depending on the operation.

Language Definition

The following table provides language definitions.

Node Name	Attribute Name	Description
and		Performs bitwise AND between two numbers.
	name	Variable to AND result into.
	value	Number to AND into result.
or		Performs bitwise OR between two numbers.
	name	Variable to OR result into.
	value	Number to OR into result.
increment		Performs ADDITION of two numbers.
	name	Variable containing the initial value AND to receive ADDITION results.
	value	Number to ADD to initial value.
decrement		Performs SUBTRACTION of two numbers.
	name	Variable containing initial value AND to receive SUBTRACTION results.
	value	Number to SUBTRACT from initial value.
divide		Performs DIVISION of two numbers.
	name	Variable containing the initial value AND to receive DIVISION results.
	value	Number by which to divide the initial value. Division by zero generates an error and stops any further processing of the current session by this parser.
modulo		Performs MODULO of two numbers.
	name	Variable containing the initial value AND to receive MODULO results.

Node Name	Attribute Name	Description
	value	Number by which to divide the initial value. Division by zero generates an error and stops any further processing of the current session by this parser.
multiply		Performs MULTIPLICATION of two numbers.
	name	Variable containing the initial value AND to receive MULTIPLICATION results.
	value	Number by which to MULTIPLY the initial value.
shiftleft		Performs a binary shift left.
	name	Variable containing the initial value AND to receive shift results.
	value	Number of bits to shift by.
shiftright		Performs a binary shift right.
	name	Variable containing the initial value AND to receive shift results.
	value	Number of bits to shift by.

Common Parser Operations

This topic provides some examples of common parser operations.

This topic includes five common parser operations.

Match Port and Identify Immediately

```
<?xml version="1.0" encoding="utf-8"?>
<parsers
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="parsers.xsd">
  <parser name="CustApp" desc="Acme Custom App" service="45324">
    <declaration>
      <port name="port" value="45324" />
    <declaration>
      </match name="port">
        <identify />
      </match>
    </parser>
  </parsers>
```

Match Port and Delay Identification

```
<?xml version="1.0" encoding="utf-8"?>
<parsers
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="parsers.xsd">
  <parser name="MSRPC" desc="Microsoft RPC protocol" service="135">
    <declaration>
      <port name="port" value="135" />
      <number name="state" scope="session" />
      <session name="end" value="end" />
    </declaration>
    <match name="port">
      <assign name="state" value="1" />
    </match>
    <match name="end">
      <if name="state" equal="1" />
        <identify />
      </if>
    </match>
  </parser>
```

```
</parsers>
```

Match Token and Identify Immediately

```
<?xml version="1.0" encoding="utf-8?>
<parsers
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="parsers.xsd">
  <parser name="RDP" desc="Remote Desktop Protocol" service="3389">
    <declaration>
      <token name="signature" value="Cookie: msthash=" />
    </declaration>
    <match name="signature">
      <identify />
    </match>
  </parser>
</parsers>
```

Match Multiple Tokens

```
<?xml version="1.0" encoding="utf-8"?>
<parsers
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="parsers.xsd">
  <parser name="MyServiceMultiToken" desc="Multiple Tokens" service="333">
    <declaration>
      <number name="state" scope="stream" />
      <token name="user" value="USER " />
      <token name="pass" value="PASS " />
      <session name="session" value="end" />
    </declaration>
    <match name="user">
      <or name="state" value="1" />
    </match>
    <match name="pass">
      <or name="state" value="2" />
    </match>
    <match name="session">
      <if name="state" equal="3">
        <identify />
      </if>
    </match>
  </parser>
</parsers>
```

```
    </parser>
</parsers>
```

Match Token and Create Metadata

```
<?xml version="1.0" encoding="utf-8"?>
<parsers xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="parsers.xsd">
    <parser name="SHELL" desc="Command Shell Identification">
        <declaration>
            <token name="cmd.exe" value=" (C) Copyright 1985-2001 Microsoft
Corp" options="linestart" />
            <meta name="client" key="client" format="Text" />
        </declaration>
        <match name="cmd.exe"
            <register name="client" value="MS Command Shell" />
        </match>
    </parser>
</parsers>
```


General Functions

This topic defines language for the flex parser general functions.

General Functions Language Definition

Node Name	Attribute Name	Description
apptype		Gets the currently defined service type for the current session.
	name	A number variable to receive the current service type.
identify		Marks the session with the parser's service type if the service type has not already been identified.
assign		Assigns a value to a variable.
	name	The unique identifier assigned to the item in the declaration section.
	value	Optional. If specified, the action defined in the match is only applied when the declaration matches the given value.
getmeta		Retrieves the value of meta that generated a callback. This function will return empty results (0, zero length string) if called when there was no meta callback.
	name	The variable to receive the value of the meta key that generated the callback.
gettoken		Returns the current matched token.
	name	A string variable to receive the current matched token. If there is no current token, the variable is assigned an empty string.
end		This terminates the execution of the current match section.
if		Compares two values. If the comparison is true, executes any sub-actions. Comparisons can be number or string types, as long as both values are the same type.
	name	The unique variable identifier assigned to the item in the declaration section.
	equal notequal less lessequal greater greaterequal and or	The operation value to compare. If true, any sub-actions are executed.
register		Adds metadata to the session.

Node Name	Attribute Name	Description
	name	The unique identifier of a meta variable to be created, as defined in the declaration section.
	value	The value of the metadata to be created.
while		Compares two values and executes any sub-actions if the comparison is true. Comparisons can be number or string types, as long as both values are the same type.
	name	The unique variable identifier assigned to the item in the declaration section.
	equal notequal less lessequal greater greaterequal and or	Specifies the operation value to compare. If true, any sub-action is executed. The and and or attributes signify bitwise operations and can only be applied to number variables.
call		Executes the specified match element. This can be any match element defined in the same flex parser regardless of how it was declared.
	value	The name of the match element, or a string variable containing the name of a match element. • If the match element name is specified, the parser will not load if the named matched element doesn't exist. • If a string variable is specified, the call element will execute any child elements that it may have if the string value resolves to a match element after executing the named match element. • If no match element can be found matching the string value, no action is taken.

Logging Functions

This topic defines language for the flex parser logging functions.

Logging functions provide a means for a flex parser to write to the system log. Logging functions can be extremely useful when creating a new flex parser, but should be kept to an absolute minimum when a flex parser is deployed to a production system.

Language Definition

Node Name	Attribute Name	Description
failure		Logs a message to the system log with the log level Failure .
	value	A string to include as the log message.
warning		Logs a message to the system log with the log level Warning .
	value	A string to include as the log message.
info		Logs a message to the system log with the log level Info .
	value	A string to include as the log message.
debug		Logs a message to the system log with the log level Debug .
	value	A string to include as the log message.

Nodes

This topic defines language for the flex parser nodes.

Nodes Language Definition

Node Name	Attribute Name	Description
<code>parsers</code>		The root node in each definition file.
	<code>xmins:xsi</code>	Defines the namespace to use for the schema inclusion. This attribute is not required; however, language definition is not possible without it. This node must have the following value: http://www.w3.org/2001/XMLSchema-instance
	<code>xsi:noNamespaceSchemaLocation</code>	Defines the XSD schema validation file used to validate the language definition. This attribute is not required; however, language definition is not possible without it. This node must have the following value: <code>parsers.xsd</code>
<code>parser</code>		The node that defines a single parser definition. This node must be directly under the <code>parsers</code> node. There can be more than one per file.
	<code>name</code>	The name that uniquely identifies the parser. This name should be short and succinct. This is used by the system to allow enabling and disabling. It should contain only the letters [a-z] and [A-Z].
	<code>desc</code>	Provides a friendly description of what the parser does.
	<code>service</code>	The unique number assigned to the session when identified.
<code>declaration</code>		Delineates the definition. Each of these definitions can have an associated <code>match</code> entry.

Node Name	Attribute Name	Description
token		Specifies a definition for identifying a token somewhere in the session protocol. This defines a <code>match</code> callback when the specified tokens are encountered in a session payload. The <code>read</code> position is set to the byte immediately following the matched token.
	name	This is a unique identifier for the declaration.
	value	This is the exact token value to be identified.
	options	Options specify that the token should start on a new line or at an end of a line (<code>linestart</code> or <code>linestop</code>).
meta-callback		Registers a callback for the flex parser whenever meta of a specific format is created. This can be further qualified to generate callbacks only for sessions that have been identified as a specific <code>apptype</code> (for example, 80 for HTTP).
	name	Name of the match element to be executed when a callback occurs. (String)
	key	Name of the meta key that generates callbacks. (String)
	format	The data type of the meta key that will generate the meta.
	apptype	The meta callback is only generated if the session being parsed has been identified with the specified <code>apptype</code> . (Unsigned Integer, Optional)
number		Defines a numeric variable that can be referenced elsewhere within the parser definition. All numeric values are 64-bit unsigned values.
	name	This is a unique identifier for the declaration.

Node Name	Attribute Name	Description
	scope (optional)	Specifies when to reset the variable. This can either be for each side of a two-sided session or only after a new session is detected. The possible values are global , constant , stream , and session (default).
string		Defines a numeric variable that can be referenced elsewhere within the parser definition.
	name	This is a unique identifier for the declaration.
	scope (optional)	Specifies when to reset the variable. This can either be for each side of a two-sided session or only after a new session is detected. The possible values are global , constant , stream , and session (default).
port		Defines a match callback when a session is encountered using the specified port. The read position is set to the first byte of the first stream (client) in the session.
	name	This is a unique identifier for the declaration.
	value	This is the port number to identify.
session		Defines a match callback for session begin/end events. These events only occur if a token for the parser is encountered in the session.
	name	This is a unique identifier for the declaration.
	value	Specifies that processing takes place at the beginning of a new session or at the end of a session (begin or end).
stream		Defines a match callback for stream begin/end events. These events only occur if a token for the parser is encountered in the stream.
	name	This is a unique identifier for the declaration

Node Name	Attribute Name	Description
	value	Specifies that processing takes place at the beginning or at the end of a stream (begin or end).
function		Defines a <code>match</code> section that can be used as a generic function. No callbacks are associated with this declaration.
	name	This is a unique identifier for the declaration.
meta		Defines the type of data that the parser will create.
	key	Specifies the key name. The key needs to be 1-16 bytes in size.
	format	Specifies the variant type (for example, Text , IPv4 , UInt32). Refer to the SDK documentation for a full list.
pattern		Defines a regular expression variable for use by the <code>regex</code> function
	name	This is a unique identifier for the declaration.
	scope (optional)	Specifies when to reset the variable. This can be for each side of a two-sided session or only after a new session is detected. Possible values are global , constant , stream , and <code>session</code> (default).
	value (optional)	Specifies a regular expression to assign to the pattern variable. This attribute is only valid when the scope attribute is set to <code>constant</code> .
match		The possible entries for taking an action once a match criterion has been found for a declaration. These nodes can be nested to provide deeper logic. There are several categories of execution elements (functions) that can appear as children of a match element: • General • Arithmetic • String • Payload

Payload Functions

This topic defines language for the flex parser payload functions.

These functions operate on a `read` position, set at the beginning of a `match` element.

Language Definition

Node Name	Attribute Name	Description
<code>find</code>		Searches the stream payload starting at the <code>read</code> position for a provided string value. If the value is found, the offset from the <code>read</code> position is returned. Any child elements will then execute. If not found, any child elements will not execute.
	<code>name</code>	A <code>number</code> variable to receive the offset from the <code>read</code> position where the match begins.
	<code>value</code>	A string to find.
	<code>length</code> (optional)	A limit to the length of the payload to be searched. If a limit is not provided, the remainder of the payload is searched. It is recommended to always use the smallest value possible here in order to reduce the effect on performance.
<code>install-decoder</code>		To enable tokens to match on payload data that may be fragmented or otherwise encoded. A scan decoder can be installed to preprocess a section of the payload before it is scanned for tokens. An example would be an HTTP response that uses the chunked transfer encoding with gzip content encoding. By parsing the HTTP header, the necessary type, offset, and length parameters can all be set, after which the HTTP response payload would appear to the token scanning as if neither encoding had been applied. However, this incurs significant overhead.
	<code>type</code>	The type of decoder to install. Valid options are: <code>gzip</code> , <code>deflate</code> , <code>chunked</code> , <code>chunked-gzip</code> , <code>chunked-deflate</code> .
	<code>offset</code>	Offset from the current <code>read</code> position to begin decoding.
	<code>length</code>	The maximum payload length to decode.
<code>isdecoding</code>		Tests whether an installed decoder is currently active. If so, any children of this function will execute. This function has no parameters.

Node Name	Attribute Name	Description
move		Moves the <code>read</code> position forward in the current stream by a specified number of bytes. If there is sufficient data in the stream, the <code>read</code> position is updated and any child elements will then execute. If not found, the <code>read</code> position remains unchanged and any child elements will not execute.
	value	The number of bytes to move the <code>read</code> position.
	direction (optional)	The direction to move the current read position. Can be <code>forward</code> (default) or reverse .
packetid		Returns the id of the packet for the current read position. It is possible for the result to be 0, which indicates that the packet id could not be determined.
	name	A number variable to receive the current packet id.
payload-position		Returns the current read position. This is a zero based index into the stream payload.
	name	A number variable to receive the current read position.
read		Reads a specified number of bytes starting at the <code>read</code> position into a variable. If there is sufficient data in the stream, the <code>read</code> position is updated, the data read assigned, and any child elements will then execute. If not found, the <code>read</code> position remains unchanged and any child elements will not execute.
	name	The name of a <code>string</code> or <code>number</code> variable to receive stream data. If a <code>number</code> variable is provided, the bytes read are interpreted as a single unsigned numeric value.
	length	The number of bytes to read from a stream.
	endianess (optional)	The byte ordering to use when reading into a number variable. Can be <code>big</code> (default) or <code>little</code> . The attribute is invalid when reading into a <code>string</code> variable.

Regex

This topic defines language for the flex parser regex node.

Regex searches the stream payload starting at the `read` position for matches to a provided regular expression. If matches are found, the offset from the `read` position and, optionally the matched string, is returned. Any child elements execute. If no matches are found, child elements do not execute.

Language Definition

Attribute Name	Description
name	A <code>number</code> variable to receive the offset from the <code>read</code> position where the match begins.
value	A regular expression to find.
length (optional)	A limit to the length of the payload to be searched. If a limit is not provided, the remainder of the payload is searched. It is recommended to always use the smallest value possible here in order to reduce the effect on performance.
found (optional)	The name of a <code>string</code> variable to receive a matched string.

String Functions

This topic provides language definitions for the flex parser string functions.

String Functions Language Definition

Node Name	Attribute Name	Description
append		Attaches a number or string to the end of a <code>string</code> variable.
	name	The unique identifier of a string variable to which the specified value is to be attached.
	value	A number or string to attach.
find		Searches a string for a provided string value. If it is found, the position is returned and any child elements will execute. Otherwise, child elements will not execute.
	name	A <code>number</code> variable to receive the zero-based position, where the provided value string was found in the <code>in</code> string.
	value	A string to find.
	in	A string to search.
	length (optional)	A limit to the length of the <code>in</code> string to be searched. If a limit is not provided, all of <code>in</code> will be searched.
length		Assigns the length of a string to a <code>number</code> variable.
	name	A <code>number</code> variable to receive the length of the specified string.
	value	A string value whose length is to be determined.
regex		Searches a string for matches to the provided regular expression. If a match is found, the position and, optionally, the matching string is returned. Any child elements will then execute. If not found, any child elements will not execute. Regular expression operations can adversely affect system performance.
	name	A <code>number</code> variable to receive the zero-based position, where the provided regular expression matched in the <code>in</code> string.
	value	A regular expression to be searched for.

Node Name	Attribute Name	Description
	<code>in</code>	A string to search.
	<code>length (optional)</code>	A limit to the length of the <code>in</code> string to be searched. If a limit is not provided, all of <code>in</code> will be searched.
	<code>found (optional)</code>	The name of a string variable to receive the matched string.
<code>substring</code>		At least one of the optional attributes <code>from</code> and <code>length</code> must be specified.
	<code>name</code>	The unique identifier of a string variable to receive the extracted value.
	<code>value</code>	A string value from which to extract a substring.
	<code>from (optional)</code>	The zero-based position from which to begin the substring. If not specified, it defaults to zero.
	<code>length (optional)</code>	The number of characters to extract. If not specified, it defaults to the remaining length of the string.
<code>tolower</code>		Converts a string to all lowercase letters.
	<code>name</code>	The name of a <code>string</code> variable to process.
<code>toupper</code>		Converts a string to all uppercase letters.
	<code>name</code>	The name of a <code>string</code> variable to process.
<code>urldecode</code>		Decodes a string containing url-encoded characters.
	<code>name</code>	A string variable to receive the decoded string.
	<code>value</code>	A url-encoded string to decode.
<code>base64decode</code>		Decodes a base-64 encoded string.
	<code>name</code>	A string variable to receive the decoded string.
	<code>value</code>	A url-encoded string to decode.
<code>uudecode</code>		Decode a uuencoded string.
	<code>name</code>	A string variable to receive the decoded string.
	<code>value</code>	A uuencoded string. The header and trailing lines should not be included.
<code>quotedprintabledecode</code>		Decode a Quoted-printable encoded string.
	<code>name</code>	A string variable to receive the decoded string.
	<code>value</code>	A quoted-printable encoded string.

Node Name	Attribute Name	Description
convert-ebcdic		Convert an EBCDIC string to its ASCII equivalent.
	name	A string variable to receive the decoded string.
	value	A url-encoded string to decode.

GeoIP2 Parsers

This topic describes the GeoIP2 parser for Decoders. This parser converts IP addresses into geographic locations, such as the country name and city where the IP address is typically found.

Note: In version 11.3 and later, the native GeoIP2 parser replaces the GeoIP parser (which has been permanently removed). The GeoIP2 parser provides the same basic functionality as the GeoIP parser as well as many enhancements. For example, it converts IP addresses into geographic locations, provides the latest Maxmind GeoIP package, and supports IPv6 addresses as well as IPv4.

Available in NetWitness Platform version 11.2 or later, the GeoIP2 Parser is enabled by default for upgrades and new installations. The GeoIP2 parser provides the latest Maxmind GeoIP package and supports IPv6 addresses as well as IPv4.

To edit the GeoIP2 parser configuration:

1. Go to  (Admin) > **Services**.
2. In the **Administration services** view, select a Log Decoder or a Decoder.
3. Click the settings icon () and select **View > Config**. In the Parsers Configuration panel, select **GeoIP2** to view and update configuration options.
4. Define the IP addresses to lookup. The GeoIP2 parser enables the following IP addresses by default: `ip.src`, `ip.dst`, `ipv6.src`, and `ipv6.dst`. You can update options by using `parsers.options` to remove or add new IP addresses. For example, you can edit `parsers.options` and pass a comma-separated list of IP addresses to use as follows:

```
GeoIP2="ipaddr=ip.src,ip.dst,ipv6.src,ipv6.dst,alias.ip"
```

This adds a new IP address to lookup called `ip.addr`. However, since `alias.ip` does not end in `.src` or `.dst`, the parser will elect to place the GeoIP2 metadata generated in meta keys without a `.src` or `.dst` suffix. So, you would see country, city, and so on, after the `alias.ip` metadata.
5. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Note: The list you pass for `alias.ip` replaces the default list. So, if you pass `alias.ip=ip.src`, it generates only GeoIP2 metadata for `ip.src`, and generates no metadata for other IP addresses.

Note: `parsers.options` is used for passing options to multiple parsers. So if you add GeoIP2 to it, you should not delete any other options being passed to other parsers (like Entropy).

The following table provides the full list of metadata that the GeoIP2 parser can potentially generate and indicates which metadata is or is not enabled by default:

Enabled by Default	Not Enabled
country, country.src, country.dst	latdec, latdec.src, latdec.dst
	longdec, longdec.src, longdec.dst
domain, domain.src, domain.dst	isp, isp.src, isp.dst
org, org.src, org.dst	city, city.src, city.dst

You can enable the other metadata using the standard parser configurations.

Note: By disabling some metadata by default, the GeoIP2 parser does not work the same as the GeoIP parser (which did not, by default, disable any metadata it generated). If you need any of the disabled metadata, you need to enable them (once only) for each Decoder, after upgrading to 11.2 or later. Keep in mind that the `isp` and `org` meta keys usually produce an equivalent value to `domain`.

Lua Parsers

There are a number of Lua parsers available from Live. See [RSA Content](#) for:

- A complete list of these parsers
- Their interdependencies
- The Flex parsers that are subsumed by each Lua parser

Five common parser operations are:

- Match Port and Identify Immediately
- Match Port and Delay Identification
- Match Token and Identify Immediately
- Match Multiple Tokens
- Match Token and Create Metadata

HTTP Parsers

The HTTP parser is a native parser that is used for Decoders to parse both requests and responses in HTTP messages. In version 11.4 and later, the HTTP parser provides a decompression option.

The decompression option mimics the decompression option in the Lua HTTP parser and is controlled by the 'decompression' option for the HTTP parser. Parser options are set in the `/decoder/parsers/config/parsers.options` configuration node. To set an option on the HTTP parser, you append an `HTTP=""` clause to the `parsers.options` field so that HTTP can understand the 'decompression' option.

For example, you could add `HTTP="decompression=true"` to the parser option list to enable decompression of all HTTP compressed bodies.

You can use the following values in the 'decompression' field.

Value	Description	Example Entry in <code>parsers.options</code>
true	Decompress all bodies	<code>HTTP="decompress=true"</code>
false	Do not decompress. This is the default.	<code>HTTP="decompress=false"</code>
1	Decompress application/* content	<code>HTTP="decompress=1"</code>
2	Decompress audio/* content	<code>HTTP="decompress=2"</code>
4	Decompress font/* content	<code>HTTP="decompress=4"</code>
8	Decompress image/* content	<code>HTTP="decompress=8"</code>
16	Decompress message/* content	<code>HTTP="decompress=16"</code>
32	Decompress model/* content	<code>HTTP="decompress=32"</code>
64	Decompress text/* content	<code>HTTP="decompress=64"</code>
128	Decompress video/* content	<code>HTTP="decompress=128"</code>

The numeric values can be combined by addition to search for multiple types of content. For example, if you want to decompress application and text content, use $1 + 64 = 65$, which becomes `HTTP="decompress=65"`.

To set the decompression option:

1. Go to  (Admin) > **Services** and select a Decoder, and in the actions menu (), select **View** > **Explore**.
2. Expand **decoder** > **parsers** and select **config**.
3. In **parsers.options**, append `HTTP="decompress=<option from table>"`.

4. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Visibility into HTTP/2 Sessions

You can search for metadata items derived from headers in the HTTP/2 stream to gain visibility into HTTP/2 sessions.

To turn on header parsing for HTTP/2 sessions:

1. Go to **ADMIN > Services** and select a Decoder, and in the actions menu (), select **View > Explore**.
2. Expand **decoder > parsers** and select **config**.
3. In **parsers.options**, append `HTTP2="headers=true"`.
4. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Decoder Snort Detection

Snort is a free and open-source intrusion prevention system that uses a rule-based language to detect malicious network traffic. The NetWitness Platform Decoder offers compatibility with Snort detection rules, sometimes referred to as Snort signatures. NetWitness Platform supports importing existing Snort rules into the Decoder, extending the known threats that can be detected. There are many sources for Snort rules, as it is an open source solution. In most cases, the behavior of Snort rules in the Decoder is the same as the rules in other network analysis tools. However, the Decoder does not have coverage for all Snort rule options. The topics in this section highlight NetWitness Platform Decoder-specific details of how to set up Snort parsers, how Snort rules are handled, and explain cases where Decoder might differ in behavior from Snort.

Using Snort Rules in Decoders

The Snort Detection engine is implemented as a built-in native parser in Decoders. The parser's identifier is `Snort`, and can be found in the Decoder Parsers Configuration page along with the other native parsers enabled by default. The Snort parser reads files from the directory `/etc/netwitness/ng/parsers/snort`. All files in this directory are read during the Decoder service restart. Files that end in the extension `.conf` are treated as configuration files. The format for configuration files matches the configuration file format for Snort, although not all configuration directives are supported. Further details are in [Configuration Directives](#). Files that end in the extension `.rules` are treated as Snort rules files. A rules file can contain one or more rules, and the Snort directory can contain more than one rules file. The rule file format is the same as for Snort itself. The Snort rule files can be loaded into the Decoder by any of the following methods:

1. The Decoder service starts.
2. The `reload` message is sent to the `/decoder/parsers` folder. This allows for rules to be reloaded while capture is running.

However, note the following:

- Any rule that does not properly parse is ignored.
- Any valid Snort rule should successfully parse; however, there are rule options that are not supported by Decoders which are not fully parsed. For more information, see [Snort Parser Capabilities](#).

Note: RSA recommends that the maximum number of Snort rules in use per Decoder is 1,000.

Initial Decoder Snort Configuration

Follow these steps to set up an initial configuration of Snort on a Decoder.

1. Edit the `snort.conf` file in the `/etc/netwitness/ng/parsers/snort` directory, and configure it according to the instructions in [Configuration Directives](#), which describes the details of the variables inside the configuration file. If the variables `HOME_NET`, `EXTERNAL_NET` and `HTTP_PORTS` are not defined, they default to the value of `any`.

A sample `snort.conf` file is shown below, but depending on where the Decoder is placed in the network, and which threats need to be captured, you might need to alter the configuration.

For example, you can be more specific about the internal subnets by defining `ipvar HOME_NET [192.168.0.0/16, 10.0.0.0/8, 172.16.0.0/12]`. However, if the Decoder is capturing

traffic on a DMZ, those subnets might not be accurate.

Also, if `HOME_NET` has been defined, and the `EXTERNAL_NET` has also been narrowed down to `ipvar EXTERNAL_NET !$HOME_NET`, then you may miss lateral movement traffic (for example, 192.168/16 to 192.168/16) within your environment. In general we suggest you align this configuration file with how you configure the `traffic_flow` configuration for the corresponding Traffic Flow Lua parser.

An example `snort.conf` file that is a good starting point:

```
#Configure the network addresses you are protecting
ipvar HOME_NET any
#Configure the network addresses external to your environment
ipvar EXTERNAL_NET any
#Configure any specific network ports
portvar HTTP_PORTS any
```

2. Generate or upload a Snort rules file into the same directory (`/etc/netwitness/ng/parsers/snort`). Read the Snort documentation at <https://www.snort.org/#documents> for details on how to construct rules. There are many free community-generated Snort rules available as well, that can be uploaded from <https://www.snort.org/downloads/#rule-downloads>. Several other sites publish Snort signatures as responses to new threats that are found. The basic rule structure and supported sections are mentioned in [Rule Sections](#).

Note: Only the Snort V2.x rules are supported. There are some rule options that are not supported, which are described in [Snort Parser Capabilities](#).

3. Make sure the Snort configuration and rules files do not cause any errors when initialized by the Decoder by running the service restart command on the Decoder.
The reload will check in the Snort rules without restarting, but does not log any messages about whether there are any issues with the rules or configuration.
 - a. In the NetWitness Platform user interface, go to  (Admin) > **Services** > **Decoder** >  > **View** > **System**.
 - b. Select **Shutdown Service** and confirm. (The Decoder will automatically restart.)
 - c. Go to > **Services** > **Decoder** >  > **View** > **Logs**.
 - d. Search for Snort to validate that the rules files loaded.
The log will indicate if the rules files are loaded completely (`full`), partially (`partial`), or if any rules failed to load (`failed`).
The following are example log entries, using `mysnort.rules`, with two fully supported rules in the file. If any of the rules are marked partial or failure, examine the rules to determine why the load was not successful.

```
Loaded mysnort.rules, full 2, partial 0, failures 0
Loaded 2 snort rules, 2 small tokens, 0 with pcres, 0 partial, 0
unsupported
```


When you search for Snort, you can find details on configuration variables and warning messages

if they are not defined. The following is an example log entry for a `snort.conf` file that is not fully populated with the minimal IP and port variables.

```
Undefined snort ip variables (will default to any): EXTERNAL_NET, HOME_
NET
```

```
Undefined snort port variables (will default to any): HTTPS_PORTS
```

- e. Alternatively, you can perform the service restart and review of the logs from the Decoder console:

```
systemctl restart nwdecoder
tail -f /var/log/messages | grep -i snort
```

4. Import or capture network traffic that is known to trigger the Snort rules to validate they are functioning.
5. Validate that the Snort rules are getting triggered by the network traffic or by sample uploaded PCAPs. The Snort statistics can be used to track which rules have been evaluated, compared to generated metadata. For more information, see [Rule Statistics](#). The default metadata generated by Snort rules are: `sig.id`, `sig.name`, `threat.category`, `threat.source`, and `risk.num`. The mapping of Snort rule information to NetWitness Platform metadata, and how to change it, is described in [Snort Parser Output](#).

Snort Parser Output

When a Snort rule matches a session, it produces meta items. The meta items generated may change depending on the configuration of the Snort parser. Snort parser meta key usage has been updated with a new option for the Snort parser. The option, `Snort="udm=true"` (set to true by default), uses the Aligned Unified Data Model (UDM) key set. For information about UDM, see <https://community.rsa.com/community/products/netwitness/rsa-content/udm>. By default, the Aligned UDM key set is used.

Note: To pass options to parsers, you must first give the name of the parser and then the options to be passed in this format:

```
<ParserName>=<ParserOptions><Whitespace><ParserName2>=<Parser2Options>"
```

Each `ParserName=Value` option must be separated by whitespace. Normally, the Value must have double quotes around it. The Value itself can sometimes list multiple `Option=Value` pairs, each separated by whitespace, and if those values have whitespace, they must be in escaped double quotes. To escape a quote, place a backslash before it: `\`.

This is an example of defining options for Parser1, Parser2, and Parser3:

```
Parser1="Option1=\"Option1 Value With Space\" Option2=Option2ValueNoSpace"
Parser2="Option1=Value" Parser3="op1=val1 op2=val2 op3=\"another value\""
```

The following keys are generated in the Aligned Unified Data Mode:

- `sig.id` with the value of the rule's `sid` field
- `threat.category` with the value of the rule's `classtype` field
- `sig.name` with the value of the rule's `msg` field
- `risk.num` with the value of the rule's priority

The following keys are generated in Legacy Key Mode:

- `alert.id` with the value of rule's `sid` field
- `threat.category` with the value of the rule's `classtype` field
- `risk.info`, `risk.suspicious`, or `risk.warning` depending on the priority of the rule. The value of the risk meta is the rule's `msg` field.

If you want to use the legacy key set which contains keys that are consistent with previous releases for Snort parser meta keys, follow these steps.

1. In the NetWitness Platform User Interface, go to **ADMIN > Services**.
2. Select a Decoder and then click  > **View > Explore**.
3. In the left panel, select **decoder > parsers > config**.
4. In the right panel, in `parser.options`, add `Snort="udm=false"`.

Rule Statistics

The Snort parser maintains counters to monitor the activity of each rule. These counters can be retrieved using the running `snrtStat` on the Decoder's `/decoder/parsers` node, which can be accessed through the NwConsole utility or the REST interface. The statistics reported for each rule are:

- `evaluations`: This counts the number of times the rule was evaluated. A rule can be evaluated once for each stream in a session. Typically a rule will only match one side of a session, and will be evaluated, at most, once. The `content` strings determine if a rule will be evaluated in a session. A `content` pattern must match somewhere in the session in order for it to be evaluated.
- `hits`: This counts the number of times a rule matched and the corresponding alert meta item was created.

To review the Snort statistics when using the Decoder console:

1. Log in to the console locally or remotely
2. Type the command `NwConsole`
3. Log in to `localhost:50004` `<admin>` `<netwitness>` (change login credentials accordingly)
4. Enter `/decoder/parsers snrtStat`

The output is a JSON structure listing the number of evaluations and hits performed, along with the options evaluated, for each Snort ID.

Configuration Directives

Decoder supports a limited number of configuration directives.

Variable Definitions

Decoder supports definitions of variables that hold the value of IP ranges using `ipvar` or Port ranges using `portvar`. For example, Snort rules usually make use of `HOME_NET`, `EXTERNAL_NET` and `HTTP_PORTS` variables. These have to be defined in your configuration file. If they are undefined, they default to the special value `any`.

ruletype

The definition of additional rule types is supported. However, only rules that have a base rule type of `alert` are supported. Rule type definitions follow the same definition format as Snort configuration.

Other Configuration Entries

config detection: debug

Adding this configuration line turns on debug messages that are generated as the rule is being processed. These debug messages are generally emitted only emitted when a rule matches a session by one of the content patterns, but then fails to match other options within the rule. This is useful for debugging complex filters in Snort rules, to determine why they may not be matching as expected.

config nopcre

Adding this configuration line turns off the PCRE (regular expression) capability.

config classification

Classification configuration influences how meta is generated when a rule matches a session. Meta output from Snort rule matches is described in [Snort Parser Output](#).

Snort Parser Capabilities

The general format of a Snort rule is that it contains a rule header and rule options. The rule header consists of the rule action, protocol type, source-destination criteria (IP addresses, ports), and the direction of traffic. The rule options consist of the message (`msg`) which describe what the rule has detected, the references related to the threat or where the rule was generated, the classification of the rule, the unique signature (`rule`) identifier, and a long list of attributes (for example, `flow`, `content`, `pcre`) to define how the traffic is identified. The specific details on which header and rule options are supported are outlined in the remainder of this section.

Rule Sections

Section	Description
Header	The header conditions are evaluated when a rule receives the first token callback for a stream. The header is evaluated once per stream, and prevents any further consideration of a rule against a specific stream if the conditions are not met.
Actions	The specified action or a rule must be defined (either one of the native Snort actions, or defined in the configuration using the <code>ruletype</code> statement) for the rule to be considered valid. The Decoder only uses rules with alert actions.
Protocols	The Decoder supports the current Snort protocol keywords (<code>tcp</code> , <code>udp</code> , <code>icmp</code> , <code>ip</code>).
IP Addresses	The full language for defining IP addresses is supported, including lists, CIDR, and negation.
Port Numbers	The full language for defining port numbers is supported, including lists, ranges and negation.

Section	Description
Direction Operator	The directional operator supports the from-to ('->') and bidirectional ('<>') values. The to-from ('<-') value is invalid and causes the rule to fail to load.

Content

The Snort `content` statement makes up the bulk of the Snort pattern detection capability. For best performance, Decoder strives to map `content` patterns to the Token Parser used by most NetWitness parsers. Therefore, the `content` strings are the main mechanism to activate the rule matching engine on a session for a particular rule. If the rule does not have any content patterns, it is effectively unsupported. The Snort parser allows for very short strings patterns in content, but be aware that a rule that contains only very short content may have to be evaluated for every session, which will make the Snort Engine run much more slowly.

The Snort parser supports most forms of `content` statements, including the following modifiers. All modifiers effect the last `content` pattern declared in the rule.

Option	Description
<code>nocase</code>	Case-insensitive pattern matching is used.
<code>offset</code>	This option is applied to the distance of the token from the beginning of the packet. If the number of bytes between the offset point and the start of the token is less than this value, it is not a match.
<code>depth</code>	This option is applied to the distance of the token from the current offset position, or the beginning of the packet if no offset is set. If the number of bytes between the 'offset' point and the beginning of the token is greater than this value, it is not a match.
<code>distance</code>	This option is applied to the distance of the token from the current Detection Offset End (DOE) point. If the number of bytes between the current token and the DOE is less than this value, the token is not a match.
<code>within</code>	This option is applied to the distance of the token from the current Detection Offset End (DOE) point. If the number of bytes between the current token and the DOE is greater than this value, the token is not a match.

These content modifiers are supported if the Snort parser is used in conjunction with AppHTTP native parser:

Option	Description
<code>http_client_body</code>	The token must appear in the body of an HTTP request.
<code>http_cookie</code>	The token must appear in an HTTP cookie header.
<code>http_header</code>	The token must appear any where in in an HTTP request or response header.

Option	Description
<code>http_method</code>	The token must appear in the HTTP method field. The detection engine does not currently restrict the method tokens, such as <code>GET</code> or <code>POST</code> , to the HTTP method field, so these matches will work even if <code>AppHTTP</code> is turned off.
<code>http_uri</code>	The token must appear in the normalized HTTP header. Non-normalized header matches are not supported.

Further detail on the HTTP related modifiers is described below in [HTTP Parsing](#).

Flow

The flow directive verifies that the rule is only applied to the client or server stream. It supports the following options.

Option	Description
<code>to_client</code>	Limits the rule to only matching on a stream that a Decoder has defined as Server.
<code>from_server</code>	Synonym for <code>to_client</code> .
<code>from_client</code>	Limits the rule to only matching on a stream that a Decoder has defined as Client.

Fast Pattern

The Snort parser uses content tagged with `fast_pattern` to indicate tokens that should always be used as tokens in the Decoder's token scanner. Decoder supports the `fast_pattern` and `fast_pattern:only` directives, but it does not support the `fast_pattern:offset,length` directive.

PCRE

The Snort parser implements `pcre` regex matching using the same `libpcre` library as other tools. It supports most of the flags documented for the `pcre` tag, with a few minor exceptions as described later in this document.

HTTP Parsing

The Snort parser relies on the native HTTP parser to notify it when HTTP is being parsed. The HTTP parser feeds information to the Snort engine to let it know which regions of the session should be used to evaluate the `http content` flags. The `pcre` implementation has option flags analogous to the `http content` options, and these are supported as well. If the native HTTP parser is disabled, then these features are also disabled in the Snort parser.

File Data

The Snort parser supports the `file_data` tag for HTTP response bodies only. This functionality relies on the native HTTP parser being enabled. If the Native HTTP parser has the decompression feature turned on, the Snort parser also uses decompressed responses to implement rule matching inside of `file_data` regions. You can enable the HTTP decompression feature by adding the token `HTTP="decompress=true"` to the configuration field `/decoder/config/parsers.options`.

The Decoder has a Javascript normalization parser that functions similarly to the `normalize_javascript` feature in Snort's HTTP preprocessor. It also affects data presented to Snort rules using `file_data`. To enable the Javascript normalization parser, add the token `JSNormalize="unescape=true"` to the configuration field `/decoder/config/parsers.options`.

The Snort parser emulates the `file_data` behavior in the same way as Snort. Directives that appear after `file_data` are implicitly evaluated whenever a content match is found inside of a `file_data` region. Because of this, `file_data` has side effects on other statements that appear after it in a rule.

Binary (byte) Directives

The Snort parser supports the `byte_test`, `byte_extract` and `byte_jump` directives.

The detection engine supports most options on these directives:

- The options for extracting the value, such as the offset, the little endian / big endian options, and relative options are supported. The relative flag causes the extraction of the byte value to occur relative to the DOE (Detection Offset End) pointer.
- `byte_jump` and `byte_extract` support the `align` and `multiplier` options to manipulate the extracted value.

The `'byte_extract` directive stores values in variable names that can be used to substitute a number into a few specific locations elsewhere in the rule. These locations are supported:

- The `value` field in a subsequent `byte_test`
- The value of a `within` statement on a subsequent content pattern
- The value of a `distance` statement on a subsequent content declaration.

The `byte_jump` option explicitly moves the DOE (Detection Offset End) pointer for use by subsequent directives, like `byte_test`, content marked with `distance` or `within`, or PCRE patterns marked as `relative`.

There are some less common byte operations that are not supported:

- The byte options do not support the `string` option for converting text into a binary value
- The byte options do not support the `bitmask` option.
- The byte options do not support the `dce` option.

Default option / fallback

When the Snort parser encounters rule options it does not yet support, it skips those options and assumes they did not prevent the rule from matching. In such cases it relies on the implementation of options it does support, such as the content or pcre statements, to determine if the rule matches. This behavior means that the Snort parser is more likely to fail in the `false-positive` direction rather than the `false-negative` direction when using features that Decoder does not yet support.

Packet Scanning/Scoping

Depending on how the rule is formed, it may be intended to match content entirely within single packets, or it may be intended to match data that spans multiple packets. The Snort parser attempts to handle both cases. Rule statements that are implicitly packet-scoped may be evaluated for every packet available in a session at parse time, and the Snort parser correctly scopes such rules so they only match if the elements of the rule match in a single packet. Conversely, if the rule uses a feature that implies reassembly and reconstruction, such as `file_data`, the Snort parser will allow the rule to be evaluated for the reconstructed portion of the stream that the rule references. This allows the Snort parser to emulate the various buffers that the Snort product provides.

Known Limitations for Snort Parser

Decoder Token Match Limits

The Snort Parser is activated by token scanning in Decoder, and therefore is subject to the same limitations on matches as all Decoder parsers. Configuration items such as `parse.bytes.min`, `parse.bytes.max`, `timeouts`, and `size limits` all apply to the Snort parser. The Decoder configuration guide describes these parameters in great detail. The Snort parser is implemented on top of reassembled sessions only so that features that require an assembled session can be supported.

Flow Bits

Flow bit directives, which start with the keyword `flowbits`, are currently mostly ignored by the Snort detection engine. Statements in rules that test flow bits currently always pass. However, the detection engine will honor flowbit directives that specify `noalert`. Rules marked as `noalert` never generate meta, even if they match.

DCE Preprocessor

The Snort parser does not currently implement any DCE preprocessors, so any features that rely on the DCE preprocessor are not implemented.

Thresholds and Filters

The Snort parser implements the `threshold filters` option only. It does not yet implement the `detection_filter` filter.

Raw Byte Scanning

The Snort parser does not currently allow rules to individually bypass the normalization or decoding using tags like `rawbytes`.

IPv6

The Snort rule parser does not currently allow rules to utilize IPv6 addresses as part of their header.

Note: Only the Snort V2.x rules are supported.

Performance Considerations

The performance of the Snort parser is dependent entirely on the rules that are being evaluated. Therefore it will vary greatly from one installation to another.

Content Tokens

Content tokens form the main entry point into the Snort engine. Hence they are the main factor affecting rules performance. The more times that a content string is found in the packet stream, the more work the Snort parser performs, and the more CPU time it uses. Therefore, it's important to use content tokens that are long, unique, and unlikely to appear in random data. Short tokens will spuriously activate the Snort parser, wasting CPU time.

Negation

Statements like `content` and `pcrc` support negation, meaning they only allow the rule to match if a content string is not found or if a regex doesn't match. These are usually detrimental to performance because they still tend to contain strings that will activate rule evaluation, yet will often fail since they are intended to filter out patterns. For example, a negated `pcrc` search still has to be evaluated if it is negated, and if it is used to filter out common text. it will produce lots of CPU activity on traffic that does not actually match the Snort rule.

Search Parser

The Search Parser is a custom parser used to generate metadata by scanning for predefined keywords and regular expressions. The parser searches the payload of a reconstructed session for string matches and can execute a regular expression search. You can configure the parser by editing the **search.ini** file.

Caution: The search parser can have a significant impact on system performance. It is important that both the search mechanism and the data to which it is applied to be well understood before creating new search definitions and enabling the search parser.

The search definition is used across all protocols. There are three basic search methods:

- **Keyword:** Search a stream for a specific set of words
- **Pattern:** Search a stream for a regular expression match
- **Keyword + Pattern:** Search a stream for a regular expression if it contains any of a given set of keywords.

Search Methods

The Search parser uses three basic search methods:

- **Keyword:** Search a stream for a specific set of words.
- **Pattern:** Search a stream for a regular expression match.
- **Keyword+Pattern:** Search a stream for a regular expression if it contains any of a given set of key words.

Syntax

```
Maxrecon=<max_size>Maxsearch=<max_ssearch_length>MatchLimit=<max_matches_per_
stream
Search Name
Services=<service_id_list>Keywords=<keyword_list>|Pattern=<expression>Case=0|1
Proximity=<number_of_bytes>Recon=0|1
Raw=0|1
```

Parameters

Parameters used in this command:

Parameter	Description
autocheck	Automatically fixes all problems without prompting
header Only	Check/display the header of each file

Parameter	Description
chatty	Displays a hex dump of every object in the file (huge amount of data)
dump#-#	Indicates a zero-based object or range of objects in the file to output in hex to the console

Example

Following is an example of the command:

To check all NetWitness database files located in the Collection named Default. If any problems are found, the command will describe the problem and ask if you would like to fix it.

```
dbcheck C:\Documents and Settings\User\My Documents\NetWitness\  
Investigations\Default\*.nw*
```

Wireless LAN Configuration

One of the files available for editing in the Services Config view > Files tab is `wlan-config.xml`, the wireless LAN configuration file.

It controls the 802.11 parsers. Its chief purpose is to control decryption of raw 802.11 frames captured by the Decoder. This file is optional. If decryption of 802.11 traffic is not desired, there is no need to create the file.

There are five link-level parsers related to wireless LAN packet capture:

- IEEE 802.11 parser (data frames and beacons only)
- Radiotap w/ 802.11 header
- Absolute Value Systems (AVS) w/ 802.11 header
- Prism II w/ 802.11 header
- CACE's "Per Packet Information" (PPI) w/ 802.11 header

The 802.11 wireless parsers introduced in 9.8 all share a single configuration file. This `wlan-config.xml` file is used to define any wireless access points the user may have in the network, and its primary purpose is to control decryption. The BSSID of the access point and the SSID that it's authoritative for is added to this file as well as all of the active default keys used by the access point.

Troubleshooting Parsers

This topic provides guidance for troubleshooting issues related to parsers.

Lua Parser Errors

Lua parsers occasionally generate errors. If a parser enters a state where it generates multiple errors, these errors can hinder performance. Beginning with 11.4, a new option is available for Lua parsers that instructs the Decoder to automatically disable the parser after a configurable number of errors.

The value is set on the `/decoder/parsers/config/parsers.options` configuration node as shown in the following example:

```
Entropy="log2=true" GeoIP2="ipaddr=ip.src,ip.dst,ipv6.src,ipv6.dst" error_parser="errorMax=5" addy_parser="errorMax=10"
```

This configuration node enables you to set options for different parsers. In this example, the Lua parser `error_parser` is configured with a maximum error count of 5, and the `addy_parser` to 10. The `errorMax` setting has a valid range of values from 0 (meaning the feature is disabled) to 65,535, and takes effect when parsers are loaded or reloaded.

To disable a Lua parser after a defined number of errors:

1. Go to **ADMIN > Services**, select a Decoder and then select   > **View > Explore**.
2. In the left panel, expand **decoder > parsers**, and select **config**.
3. In **parsers.options**, add the following command, where `<any_parser>` is the Lua parser for which you want to limit errors, and `<n>` is the number of errors to which the parser is limited:
`<any_parser>="errorMax=<n>"`

4. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Results of Automatically Disabling a Parser

When a parser is automatically disabled, a log message is generated (per parser thread) and states that the parser has been disabled, as shown in the following example:

```
(W) 2019-Apr-25 16:25:33 [Parse] Lua parser 'error_parser' has been disabled because it exceeded the configured error threshold (5)
```

If detailed statistics for parsers are enabled, the text of the last error is populated in the parser's detailed statistics under `/decoder/parsers/definitions/<parser-name>`. Also, a new attribute has been added to the XML returned by the `/decoder/parsers?msg=schema` call. When the parser is in an error state, the attribute `error` is set to 1 and the `enabled` attribute is set to 0. When the parser is reset, the values are reset to 0 and 1 respectively.

Resetting the Parser

If a parser is disabled because it has exceeded the error threshold, it can be reset by reloading the parser. This causes the statistics to reset so that if the faulty parser is still in place, it will function again until the error threshold is met.

Configure Feeds

NetWitness Platform uses feeds to create metadata based on externally defined metadata values. A feed is a list of data that is compared to sessions as they are captured or processed. For each match, additional metadata is created. This data could identify and classify malicious IPs or incorporate additional information such as department and location based on internal network assignments. Some examples of feeds include threat feeds to identify BOTNets, DHCP mappings, or even Active Directory (AD) information such as physical location or logical department.

You can use the Live module in NetWitness Platform to obtain feeds from outside sources. "Live Content in NetWitness Platform" in the *Live Services Management Guide* provides an overview of the Live content management tool.

Within the NetWitness Platform user interface, you can view the list of currently deployed feeds, along with an indicator if a feed that originated from Live was installed through NetWitness Platform or manually. Feeds can be added, removed, and updated while a Decoder is running without affecting capture.

Custom Feed Definition File Structure

The NetWitness Platform Custom Feed wizard allows creation and deployment of custom Decoder feeds based on deterministic logic that offers the meta keys specific to the selected Decoders and Log Decoders. Although the wizard guides users through the process to create both on-demand and recurring feeds, it is helpful to understand the form and content of a feed file when you create a feed.

Before you create or modify custom feeds, keys, or data types, ensure that you have a good understanding of the relationship between keys and datatypes in databases. It is important that any custom keys that you create match the corresponding datatypes. If your feed utilizes a custom meta key, you must define that key's data type in `index-decoder-custom.xml`. This is to ensure that the key data is generated with correct data format.

Also, to enable investigation on the custom meta key, you must define its index on the concentrator, using the `index-concentrator-custom.xml` file. For more information on how to write the `index-decoder-custom.xml` or `index-concentrator-custom.xml` file, see the topic "[Index Customization](#)" in the *Core Database Tuning Guide for RSA NetWitness Platform*.

Feed filenames in RSA NetWitness Platform are in the form `<filename>.feed`. To create a feed, NetWitness Platform requires a feed data file in `.csv` or `.xml` format and a feed definition file in `.xml` format, which describes the structure of a feed data file. The Custom Feed wizard can create the feed definition file based on a feed data file, or based on a feed data file and the corresponding feed definition file.

The files that you use to create an on-demand feed must be stored on your local file system. The files used to create a recurring feed must be stored at an accessible URL, whence NetWitness Platform can fetch the most current version of the file for each recurrence. After a NetWitness Platform feed is created, you can download the feed to your local file system, edit the feed files, and then edit the NetWitness Platform feed to use the updated feed files.

Sample Feed Definition File

This is an example of a feed definition file named `dynamic_dns.xml`, which NetWitness Platform creates based on your entries in the Custom Feed wizard. It defines the structure of the feed data file named `dynamic_dns.csv`.

Note: The feed file path should be `.csv` regardless of the Feed Type (Default or STIX).

```
<?xml version="1.0" encoding="utf-8"?>
  <FDF xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:noNamespaceSchemaLocation="feed-definitions.xsd">

    <FlatFileFeed name="Dynamic DNS Domain Feed"
      path="dynamic_dns.csv"
      separator=","
      comment="#"
      version="1">

      <MetaCallback
        name="alias.host"
        datatype="Text"
        apptype="0"
        truncdomain="true"/>

      <LanguageKeys>
        <LanguageKey name="threat.source" datatype="Text" />
        <LanguageKey name="threat.category" datatype="Text" />
      </LanguageKeys>
    </FlatFileFeed>
  </FDF>
```

```
        <LanguageKey name="threat.desc" valuetype="Text" />
    </LanguageKeys>

    <Fields>
    <Field index="1" type="index" key="alias.host" />
    <Field index="4" type="value" key="threat.desc" />
    <Field index="2" type="value" key="threat.source" />
    <Field index="3" type="value" key="threat.category" />
    </Fields>
</FlatFileFeed>

</FDF>
```

Define Multiple Values in a Single Field

The information in this section applies to RSA NetWitness Platform version 11.3.1 and later.

Custom feeds support multiple values in a single field. This allows feeds to generate multiple values of the same meta.

For example:

```
<Fields>

    <Field index="1" type="index" key="lc.cid" />
    <Field index="2" type="value" key="user" separator=";" openchar="("
    closechar=")" />

</Fields>
```

The second field is defined to accept multiple values in the corresponding CSV file. The following is an example of the defined values:

```
clcl1,(bob;tom;sam) clcl2,susan clcl23,doris
```

This definition means that when we see the `lc.cid` with a value of `cic1`, three username meta values are generated with the values of `bob`, `tom` and `sam`.

The separator in the field definition for a multi-value field **must** be different than the separator for the feed file itself. Escaping the separator character in field values is **not** supported.

Feed Definition Equivalents for Custom Feed Wizard Parameters

The NetWitness Platform Custom Feed wizard provides options to define the structure of the data feed file. These correspond directly to attributes in the feed definition (`.xml`) file.

NetWitness Platform Parameter	Feed Definition File Equivalent
(Define Feed Tab) Feed Type	Select: Default - to define a feed based on a <code>.csv</code> formatted feed data file. STIX - to define a feed based on STIX formatted <code>.xml</code> file.

NetWitness Platform Parameter	Feed Definition File Equivalent
(Define Feed Tab) Feed Task Type	Select: Adhoc - to create an on-demand feed. Recurring - to update the .csv or .xml file persistently and store it in a location accessible by NetWitness Platform, so NetWitness Platform downloads a file at regular intervals and pushes it to the downstream devices.
(Define Feed tab) Name	The custom feed name in the feed data file. It corresponds to the <code>flatfeedfile name</code> attribute in the feed definition file. For example, Dynamic DNS Test Feed. Note: You can use special characters to define the name of the custom feed.
(Define Feed tab) File/Browse	This is the name of the feed data file. It corresponds to the <code>flatfeedfile path</code> attribute in the feed definition file. For example, <code>dynamic_dns.csv</code> .
(Advanced Options tab) XML Feed File	The name of the feed definition file. For example, <code>dynamic_dns.xml</code> .
(Advanced Options tab) Separator	The separator character used to separate attributes in the feed data file. It corresponds to the <code>flatfeedfile separator</code> in the feed definition file. For example, a comma.
(Advanced Options tab) Comment	The character used to identify a comment in the feed data file. It corresponds to the <code>flatfeedfile comment</code> attribute in the feed definition file. For example, #.
(Define Columns tab, Define Index) Type	The type of lookup value in the index position of the feed data file. IP means that each row in the feed data file contains an IP address in the lookup value position. The IP value is in dotted-decimal format (for example, 10.5.187.42). IP Range means that each row in the feed data file contains a range of IP addresses in the lookup value position. The IP range is in CIDR format (for example, 192.168.2.0/24). Non IP means that the each row in the feed data file contains a metadata value other than IP address in the lookup value position. The Service Type and Truncate Domain, and Callback Keys fields become active for a Non IP index.
(Define Columns tab, Define Index) CIDR	Specifies that the IP value in the lookup position is in CIDR format. The CIDR attribute sets the IP address format in the field to Classless Inter-Domain Routing (CIDR) notation.
(Define Columns tab, Define Index) Service Type	For a Non IP index, the integer service type to filter meta lookups. It corresponds to the <code>MetaCallback apptype</code> attribute in the feed definition file. A value of 0 indicates no filtering by service type.

NetWitness Platform Parameter	Feed Definition File Equivalent
(Define Columns tab, Define Index) Truncate Domain	For a Non IP index, for meta values that contain domain names (for example, hostnames), the system can strip off the host specific element in the data. Truncate Domain corresponds to the <code>MetaCallback truncdomain</code> attribute. If the value is <code>www.example.com</code> , it is truncated to <code>example.com</code> . A value of False selects no truncation, and True selects truncation.
(Define Columns tab, Define Index) Callback Keys	For a Non IP index, the available meta keys to match on instead of <code>ip.src/ip.dst</code> (the defaults for IP index type) are selectable from the drop-down list. The Callback Key corresponds to the <code>MetaCallback name</code> attribute, and the index column of the <code>csv</code> file must contain data that can match the chosen meta key. For example, if the <code>user</code> meta key is chosen, the index column of the <code>.csv</code> file needs to be populated with users to be matched.
(Define Columns tab, Define Index) Index Column	Identifies the column in the feed data file that provides the lookup value for the row. Each position in each row of the feed data file is identified by a Field index attribute in the feed definition file. A field with an index of 1 is the first entry in a row, the second field has an index of 2 , the third field has an index of 3 , and so on.
(DEFINE VALUES) Key	The name of the <code>LanguageKey</code> , as defined in the feed definition file, for which meta is created from this row of the feed data file. It corresponds to the <code>Field key</code> attribute in the feed definition file. A key applies only to a field whose type is set to <code>value</code> . In the feed definition file, there is a list of <code>LanguageKeys</code> from <code>index.xml</code> , or a summary name if Source Name and Destination Name are used. For example, <code>reputation</code> is a summary name for <code>reputation.src</code> and <code>reputation.dst</code> . This value is referenced by the <code>Field key</code> attribute.

Sample Files for a MetaCallback Feed Using CIDR Index Range for IPv4 and IPv6

These sample files demonstrate how to use CIDR index ranges for IPv4 and IPv6 in custom MetaCallback feeds. As with other custom feeds, you must create feed data file in `.csv` format, and a feed definition file in `.xml` format.

Note: Using MetaCallback feeds with CIDR index ranges is supported only through the Advanced Configuration wizard or the REST interface.

The following example shows the contents of both a `.csv` file and an `.xml` file for a MetaCallback feed using CIDR index ranges for IPv4 or IPv6.

.csv file:

```
192.168.0.0/24, Sydney
192.168.1.0/24, Melbourne
```

.xml file:

```
<?xml version="1.0" encoding="UTF-8"?>
<FDF xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="feed-definitions.xsd">
<FlatFileFeed name="ip_test" path="ip_test.csv" separator="," comment="#">
```

```
<MetaCallback name="DstIP" valuetype="IPv4" apptype="0"
truncdomain="false">
  <Meta name="ip.dst"/>
</MetaCallback>
<LanguageKeys>
  <LanguageKey name="alert" valuetype="Text" />
</LanguageKeys>
<Fields>
  <Field index="1" type="index" range="cidr"/>
  <Field index="2" type="value" key="alert" />
</Fields>
</FlatFileFeed>
</FDF>
```

Note: To configure a CIDR index range for feeds with single or multiple MetaCallbacks of value type IPv4 or IPv6, the field of type index MUST contain a range attribute with `range="cidr"`. Also, configuring "cidr" index ranges for feeds with MetaCallbacks of multiple different value types is not supported.

Feed Definitions File

This topic introduces the feed definitions file, which is available for editing in the Services Config view > Files tab. One of the files available for editing in the Services Config view > Files tab is **feed-definitions.xml**, the feed definitions file.

You can define feeds in the `feed-definitions.xml` file. The Decoder uses an XML schema to define feed messages when it creates a binary `.feed` file from the feeds defined here.

For details on the feed definition language, refer to [Custom Feed Definition File Structure](#)

Create a Custom Feed

You can create a custom feed using the Custom Feed wizard. To complete this procedure, you need a feed data file in .csv or .xml format. If you also have an associated feed definition file in .xml format, which describes the structure of the feed data file, you can use the feed definition file to create a feed. The Custom Feed wizard can create the feed based on a feed data file, or based on a feed data file and corresponding feed definition file.

Before you create or modify custom feeds, keys, or data types, ensure that you have a good understanding of the relationship between keys and datatypes in databases. It is important that any custom keys that you create match the corresponding datatypes. If your feed utilizes a custom meta key, you must define that key's data type in `index-decoder-custom.xml`. This is to ensure that the key data is generated with correct data format.

Also, to enable investigation on the custom meta key, you must define its index on the concentrator, using the `index-concentrator-custom.xml` file. For more information on how to write the `index-decoder-custom.xml` or `index-concentrator-custom.xml` file, see the topic "[Index Customization](#)" in the *Core Database Tuning Guide for RSA NetWitness Platform*.

Note: For information about STIX and creating a STIX custom feed, see "Create a STIX Custom Feed" in the *Decoder and Log Decoder Configuration Guide*.

The feed data file and optionally the feed definition file (.xml) must be available on the local file system for an on-demand custom feed. For a recurring custom feed, the files must be available at a URL that is accessible to the NetWitness Platform server.

Note: When you create a source and destination-based feed on a Log Decoder, it only populates the source meta key. You cannot use a range-based or CIDR feed. You must list every single IP address. To resolve this issue, create two different feeds using IP addresses and you can use CIDR in these feeds.

To create a custom feed:

1. Go to **Configure > Custom Feeds**.

The Custom Feeds view is displayed.

Name	Trigger	Disk Usage	Created	Last Run Time	Status	Progress
TEST	Once	-	2019-07-13 13:30:36	2019-07-13 13:30:36	Completed	
TEST2	Once	-	2019-07-13 13:50:33	2019-07-13 13:50:33	Completed	
te	Once	-	2019-07-14 04:16:51	2019-07-14 04:16:51	Completed	
onlydom	Once	-	2019-07-14 04:21:37	2019-07-14 04:21:37	Completed	
PCAP	Once	-	2019-07-14 09:30:49	2019-07-14 09:30:49	Completed	

2. In the **Feeds** panel, click > **Custom Feed > Next**.

The Configure a Custom Feed wizard is displayed, with the Define Feed form open.

The screenshot shows a wizard window titled "Configure a Custom Feed" with a close button (X) in the top right corner. The wizard has four steps: "Define Feed" (active), "Select Services", "Define Columns", and "Review". The "Define Feed" step contains the following fields and controls:

- Feed Type:** Radio buttons for **CSV** (selected) and **STIX**.
- Feed Task Type:** Radio buttons for **Adhoc** (selected) and **Recurring**.
- Name ***: A text input field.
- Upload As Csv File Feed:** A checkbox that is currently unchecked.
- File ***: A text input field with the placeholder "Select File" and a **Browse** button.
- Advanced Options:** A section header with a dropdown arrow, currently collapsed.

At the bottom of the wizard are four buttons: **Reset**, **Cancel**, **Prev**, and **Next** (highlighted in blue).

3. Select the Feed Type: **CSV** or **STIX**.
4. To define a feed based on a .csv formatted feed data file, select **CSV** (which is the default) in the **Feed Type** field.
5. To define an on-demand feed task that executes once, select **Adhoc** in the **Feed Task Type** field and do one of the following:

- a. (Conditional) To define a feed based on a CsvFileFeed file, select the **Upload as Csv File Feed** checkbox, type the feed **Name**, select a .csv content file from the local file system, and click **Next**. If you do not select the checkbox, the .csv file will be a FlatFileFeed file.

Note: When you select the Upload as Csv File Feed checkbox, the XML feed options under Advanced are unavailable.

- b. (Conditional) To define a feed based on an XML feed file, select **Advanced Options**.

Note: Ensure that the Upload as Csv File Feed checkbox is deselected.

The Advanced Options are displayed:

The screenshot shows a dialog box titled "Configure a Custom Feed" with a close button (X) in the top right corner. The dialog has four tabs: "Define Feed" (active), "Select Services", "Define Columns", and "Review".

Under the "Define Feed" tab, the following options are visible:

- Feed Type:** Radio buttons for ☒ CSV and ☐ STIX.
- Feed Task Type:** Radio buttons for ☒ Adhoc and ☐ Recurring.
- Name ***: A text input field.
- Upload As Csv File Feed**: A checkbox, currently unchecked.
- File ***: A text input field with "Select File" and a "Browse" button.
- Advanced Options**: A section with a collapse icon (upward arrow) and a horizontal line.
 - XML Feed File**: A text input field with "Select File" and a "Browse" button.
 - Separator**: A text input field containing a comma (,).
 - Comment**: A text input field containing a hash symbol (#).

At the bottom of the dialog, there is a horizontal scrollbar and a row of buttons: "Reset", "Cancel", "Prev", and "Next".

- c. Select an XML feed file from the local file system, choose the separator (default is comma), specify the comment characters used in the feed data file (default is #), and click **Next**. The Select Services form is displayed. This is an example of the form for a feed based on a feed data file with no feed definition file. If you are defining a feed based on a feed definition file, the Define Columns tab is not needed.

Configure a Custom Feed

Define Feed Select Services Define Columns Review

Services Groups

☐		Name ^	Address	Type
☐				Decoder
☐				Decoder
☐				Decoder
☐				Log Decoder
☐				Decoder
☐				Log Decoder
☐				Decoder
☐				Decoder
☐				Decoder
☐				Log Decoder
☐				Log Decoder
☐				Decoder
☐				Decoder
☐				Log Decoder

Reset Cancel Prev Next

6. To define a recurring feed task that executes repeatedly at specified intervals, during a specified date range:

- a. In the Define Feed form, select **Recurring** in the **Feed Task Type** field.

The Define Feed form includes the fields for a recurring feed.

Configure a Custom Feed

Define Feed | Select Services | Define Columns | Review

Feed Type: ☒ CSV ☐ STIX

Feed Task Type: ☐ Adhoc ☒ Recurring

Name:

Upload As Csv File Feed: ☐

URL: **Verify**

☐ Authenticated

☐ Use Proxy

Recur Every:

☐ Date Range

Advanced Options

XML Feed File: **Select File** **Browse**

Separator:

Comment:

Reset **Cancel** **Prev** **Next**

- b. In the **URL** field, enter the URL where the feed data file is located, for example, `http://<hostname>/<feeddatafile>.csv`, and click **Verify**. NetWitness Platform verifies the location where the file is stored in order to enable checking for the latest file automatically before each recurrence.
- c. (Optional) If the URL has restricted access and requires authentication using your username and password, select **Authenticated**.
NetWitness Platform provides your user name and password for authentication to the URL.
- d. If you want the NetWitness server to access the Feed URL through a proxy, select **Use Proxy**. For more information on configuring a proxy, see "Configure Proxy for NetWitness Platform" in the *System Configuration Guide*. By default, the **Use Proxy** checkbox is not set.
- e. To define the interval for recurrence, do one of the following:
 - Specify the number of minutes, hours, or days between recurrences of the feed.
 - Specify recurrence every week, and select the days of the week.

- f. To define the date range for the execution of the feed to recur, specify the **Start Date** and time and the **End Date** and time.

The screenshot shows the 'Configure a Custom Feed' dialog box with the 'Define Feed' tab selected. The dialog has four tabs: 'Define Feed', 'Select Services', 'Define Columns', and 'Review'. The 'Define Feed' tab contains the following fields and options:

- Feed Task Type:** Radio buttons for 'Adhoc' and 'Recurring' (selected).
- Name *:** Text input field containing 'TestFeed'.
- URL *:** Text input field containing 'https://qasa2.netwitness.local/live/feeds'. A 'Verify' button is to the right.
- Authenticated:** Checkbox (unchecked).
- Use proxy:** Checkbox (unchecked).
- Recur Every:** A numeric input field with '3' and a dropdown menu set to 'Day(s)'.
- Date Range:** A collapsed section indicated by a downward arrow.
- Advanced Options:** A collapsed section indicated by an upward arrow. When expanded, it contains:
 - XML Feed File:** A text input field with 'Select File' and a 'Browse' button.
 - Separator:** A text input field containing a comma ','.
 - Comment:** A text input field containing a hash '#'.

At the bottom of the dialog are four buttons: 'Reset', 'Cancel', 'Prev', and 'Next'.

7. (Conditional) If you want to define a feed based on an XML feed file:
 - a. Type the feed **Name**, select **Advanced Options**. The Advanced Options fields are displayed.
 - b. Select an XML feed file from the local file system, choose the **Separator** (default is comma), specify the **Comment** characters used in the feed data file (default is #) and click **Next**. The Select Services form is displayed.

Configure a Custom Feed

Define Feed Select Services Define Columns Review

Services Groups

☐		Name ^	Address	Type
☐				Decoder
☐				Decoder
☐				Decoder
☐				Log Decoder
☐				Decoder
☐				Log Decoder
☐				Decoder
☐				Decoder
☐				Decoder
☐				Log Decoder
☐				Log Decoder
☐				Decoder
☐				Decoder
☐				Log Decoder

Reset Cancel Prev Next

8. To identify services on which to deploy the feed, do one of the following:
 - a. Select one or more Decoders and Log Decoders, and click **Next**
 - b. Click the **Groups** tab and select a group. Click **Next**.
The Define Columns form is displayed.
9. To map columns in the Define Columns form:
 - a. Define the Index type: **IP**, **IP Range**, or **Non IP**, and select the index column.
 - b. (Conditional) If the index type is **IP** or **IP Range** and the IP address is in CIDR notation, select **CIDR**.
 - c. Conditional) If the index type is **Non IP**, additional settings are displayed. Select the service type and **Callback Keys**, and optionally select the **Truncate Domain** option.

Configure a Custom Feed

Define Feed | Select Services | **Define Columns** | Review

Define Index

Type: ☐ IP ☐ IP Range ☒ Non IP

Index Column: 1 Service Type: 0 ☒ Truncate Domain

Callback Key (S):

Define Values

Column	Key
1 (Index)	
SRM_Sar	
ANCEST	

OS
access.point
accesses
action
alert
alert.id
alias.host
alias.ip
alias.ipv6
alias.mac
asn.dst
asn.src

Reset Cancel Prev Next

- d. Select the language key to apply to the data in each column from the drop-down list. The meta keys displayed in the drop-down list is based on the meta keys available for the service define values. You can also add other meta keys based on advanced expertise.

Configure a Custom Feed

Define Feed

Select Services

Define Columns

Review

Define Index

Type

☐ IP

☐ IP Range

☒ Non IP

Index Column

1

Service Type

0

☒ Truncate Domain

Callback Key (S)

action

Define Values

Column	1 (Index)	2	3	4
Key		threat.source	threat.category	threat.desc
	SRM_SaaS_ES	MXASSETInterface	AddChange	EN
	ANCESTOR	ASSETNUM	ASSETTAG	ASSETTYPE
		cent45	9164	
		cent45	9164	

Reset

Cancel

Prev

Next

- e. Click **Next**.

The Review form is displayed.

Configure a Custom Feed

Define Feed | Select Services | Define Columns | **Review**

Feed Details

Name: Testing

CSV File: AssetsImportCompleteSample.csv

Service Details

Services: Log Decoder, Decoder

Column Mapping Details

Index Type: Other

Callback Key (s): action

Truncate Domain: true

Service Type: 0

Value Columns:

Index	Value Column
1	Index
2	threat.source
3	threat.category
4	threat.desc

Reset Cancel Prev **Finish**

10. Anytime before you click **Finish**, you can:
 - Click **Cancel** to close the wizard without saving your feed definition.
 - Click **Reset** to clear the data in the wizard.
 - Click **Next** to display the next form (if not viewing the last form).
 - Click **Prev** to display the previous form (if not viewing the first form)
11. Review the feed information, and if correct, click **Finish**.
12. Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file are listed in the Feed grid and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.

RSA Investigate Respond Users Hosts Files Dashboard Reports      admin

LIVE CONTENT SUBSCRIPTIONS CAPTURE POLICIES INCIDENT RULES INCIDENT NOTIFICATIONS ESA RULES CUSTOM FEEDS LOG PARSER RULES

Feeds

☐	Name	Trigger	Feed Size	Created	Last Run Time	Status	Progress
☐	 FILEHASH	Fetches STIX feeds from 2020-May-19 03:16, running every 5 minutes	0 bytes	2020-05-19 03:25:23	2020-05-22 05:16:01	Completed	
☐	 FILESTIX	Fetches STIX feeds from 2020-May-19 03:32, running every 5 minutes	0 bytes	2020-05-19 03:32:09	2020-05-22 05:12:09	Completed	
☐	 AllIndicatorsREST	Fetches STIX feeds from 2020-May-19 04:48, running every 5 minutes	0 bytes	2020-05-19 05:03:52	2020-05-22 05:13:26	Completed	
☐	 AllIndEdited	Fetches STIX feeds from 2020-May-19 05:13, running every 5 minutes	0 bytes	2020-05-19 05:13:54	2020-05-22 05:13:54	Completed	
☐	 TAXIServer1	Fetches STIX feeds from 2020-May-19 05:44, running every 5 minutes	288 bytes	2020-05-19 05:44:38	2020-05-22 05:14:38	Completed	



RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

Create a STIX Custom Feed

Structured Threat Information Expression (STIX[™]) is a structured language for describing cyber threat information so it can be shared, stored, and analyzed in a consistent manner. For more information about STIX, see <https://stixproject.github.io/>.

You can create a custom feed using a STIX-formatted feed data file (.xml) in RSA NetWitness Platform. NetWitness Platform supports Structured Threat Information Expression (STIX) 1.0, 1.1 and 1.2 versions only.

Caution: If a STIX recurring feed is configured and you update Security Analytics from 10.6.x to NetWitness Platform 11.x, you must re-configure the STIX recurring feed.

In NetWitness Platform, STIX feeds of type Indicator or Observable that contain properties such as the IP addresses, File hashes, Domain names, URIs and Email addresses are supported. The property values in the Equals operator are supported. Attributes such as Type and Title are also read from the STIX. A STIX file with a single STIX_Package is supported.

TAXII (Trusted Automated eXchange of Indicator Information) is the main transport mechanism for cyber threat information represented in STIX. Using the TAXII services, organizations can share cyber threat information in a secure and automated manner.





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

The STIX and TAXII communities work closely together to ensure that they continue to provide a full stack for sharing threat intelligence.

Apart from the TAXII server, STIX data can also reside on a REST server and you can fetch the STIX file from the REST server by providing the URL of the REST server. For example, `http://stixrestserver.internal.com`.

The STIX feed data file and optionally the feed definition file, both in .xml format must be available on the local file system for an on-demand custom feed. For a recurring custom feed, the files must be available at a URL that is accessible to the NetWitness Platform server.

To create a STIX custom feed:

1. Go to  (Configure) > CUSTOM FEEDS.

The Custom Feeds view is displayed.





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

☐	Name	Trigger	Disk Usage	Created	Last Run Time	Status	Progress
☐	TEST	Once	-	2019-07-13 13:30:36	2019-07-13 13:30:36	Completed	
☐	TEST2	Once	-	2019-07-13 13:50:33	2019-07-13 13:50:33	Completed	
☐	te	Once	-	2019-07-14 04:16:51	2019-07-14 04:16:51	Completed	
☐	onlydom	Once	-	2019-07-14 04:21:37	2019-07-14 04:21:37	Completed	
☐	PCAP	Once	-	2019-07-14 09:30:49	2019-07-14 09:30:49	Completed	

2. In the toolbar, click .






RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

The Setup Feed dialog is displayed.



3. To select the feed type, click **Custom Feed** and **Next**.

The Configure a Custom Feed wizard is displayed, with the Define Feed dialog open.





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

Configure a Custom Feed

Define Feed

Select Services

Define Columns

Review

Feed Type

☐ CSV ☒ STIX

Name *

FILESTIX

STIX Source

All STIX Files

Recur Every

5

Minute (s)

☐ Date Range

Advanced Options

XML Feed File

Select File

Browse

Separator

~

Comment

#

Reset

Cancel

Prev

Next



RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

4. Enter the following details:
 - a. **Feed Type:** Select **STIX**, to define a feed based on a STIX formatted .xml file.
 - b. **Name:** type the feed name, to define a feed based on STIX formatted .xml file.
 - c. **STIX Source:** Select a STIX data source from the drop-down which is added in Context Hub.
 - d. **Recur Every:** Specify a recurring feed task that executes repeatedly at specified intervals.

Note: NetWitness Platform verifies the connection to the server, so that NetWitness Platform can check for the latest file automatically before each recurrence.

- e. **Date Range:** Select the checkbox and specify the date range for the feed task to recur.
5. (Optional) Select **Advanced Options**, to define a feed based on an XML feed file.
 - a. **XML Feed file:** Browse and select an XML feed file from the local file system.
 - b. **Separator:** Choose a separator (default is comma).
 - c. **Comment:** Specify the comment characters used in the feed data file (default is #).
6. Click **Next**.





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

7. The Select Services dialog is displayed. This is an example of the form for a feed based on a feed data file with no feed definition file. If you are defining a feed based on a feed definition file, the Define Columns tab is not needed.





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

Configure a Custom Feed

Define Feed

Select Services

Define Columns

Review

Services

Groups

☒		Name ^	Address	Type
☒		Log Decoder	10.1.1.1	Log Decoder
☒		Decoder	10.1.1.1	Decoder

Reset

Cancel

Prev

Next





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

8. To identify services on which to deploy the feed, do one of the following:
 - a. Select one or more Decoders and Log Decoders, and click **Next**.
 - b. In case of STIX feed, Context Hub will be selected by default and you are not allowed to deselect it. In addition, you can select one or more Decoders and Log Decoders and click **Next** or Click the **Groups** tab and select a group. Click **Next**.

If the data from the STIX server is large, the following message is displayed:





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

Configure a Custom Feed

Define Feed → **Select Services** → Define Columns → Review

Services Groups

Note : STIX content will exist in the Context Hub service by default and you are not allowed to deselect it. Select Decoders/Log Decoders/Log

Fetching sample data is taking longer than expected.
Choose one of the following options

☒		CH	Continue to Wait	Map without Sample data	
☐		LD	10.31.165.66	Log Decoder	
☒		LD85	10.31.165.85	Log Decoder	

Reset Cancel Prev **Next**





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

- If you click **Continue to Wait**, it continues to wait till the sample data is fetched or timeout (10 minutes) whichever is sooner. In case of timeout no sample data is retrieved even after 10 minutes.
- If you click **Map Without Sample data**, the mapping column is displayed without any sample data.

The Define Columns dialog is displayed.

9. To map columns in the Define Columns form:

- a. Define the Index type: **IP**, **IP Range**, or **Non IP**, and select the index column.
- b. (Optional) If the index type is **IP** or **IP Range** and the IP address is in CIDR notation, select **CIDR**.
- c. (Optional) If the index type is **Non IP**, additional settings are displayed. Select the service type and **Callback Keys**, and optionally select the **Truncate Domain** option.





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

Configure a Custom Feed

Define Feed

Select Services

Define Columns

Review

Define Index

Type

☒ IP☐ Non IP

Index Column(S)

10

☐ CIDR

Define Values

Column	1	2	3	4
Key				
Header	Indicator Title	Indicator Description	Observable Title	Observable Description
	Some Indicator	Some Indicator	domain:domain1.exa...	domain:domain1.exa...
	Some Indicator	Some Indicator	domain:domain2.exa...	domain:domain2.exa...
	indicator-domain	auto domain test	domain test	domain desc
	Another Indicator	<p>Another Indicator...	domain:domain3.exa...	domain:domain3.exa...

Reset

Cancel

Prev

Next

NETWITNESS



RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

Note:

- If the **Index Type** is Non IP, you can select multiple index columns in the **Index Column(S)**. The values from all the selected columns are merged in the first index column that you selected and the merged values are pushed to the Log Decoder for parsing. For example, in the **Index Column(S)** if you select 2,4,7 as index columns the values from the 2,4 and 7 columns are merged in the column 2 and the values are pushed to Log Decoder for parsing.
- Indexing cannot be done for the columns such as Indicator Title, Indicator Description, Observable Title, Observable Description, as the look up cannot be performed for those columns.

- d. Select the language key to apply to the data in each column from the drop-down list. The meta displayed in the drop-down list is based on the meta available for the service define values. You can also add other meta based on advanced expertise.
- e. Click **Next**.

The Review dialog is displayed.





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

Configure a Custom Feed

Define Feed

Select Services

Define Columns

Review

Feed Details

Name

FILESTIX

XML Feed File

FILESTIX-stix.xml

Service Details

Services

PH - Decoder, LH - Log Decoder

Column Mapping Details

Index Type

IP

CIDR

false

Value Columns

10

Index

24

event.desc

Reset

Cancel

Prev

Finish





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

10. Anytime before you click **Finish**, you can:

- Click **Cancel** to close the wizard without saving your feed definition.
- Click **Reset** to clear the data in the wizard.
- Click **Next** to display the next dialog (if not viewing the last form).
- Click **Prev** to display the previous dialog (if not viewing the first form)

11. Review the feed information, and if correct, click **Finish**.

Upon successful creation of the feed definition file, the Create Feed wizard closes, the feed and corresponding token file are listed in the Feed grid, and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.





RSA[®] NetWitness Platform

Version 11.6

Decoder Configuration Guide

RSA

Investigate

Respond

Users

Hosts

Files

Dashboard

Reports

Note: Health and Wellness raises alerts when the available heap memory of Context Hub server is critically low. If the status of Context Hub server is Unhealthy due to low memory. For more information on how to troubleshoot `OutOfMemoryError` on Contexthub Server, refer to "Troubleshooting" in the *Live Services Management Guide*.



Create an Identity Feed

You can create an Identity feed and populate it to selected Decoders and Log Decoders. In order to create an identity feed, you need to have:

- A Log Collector service with an Identity Feed Event Processor
- A Log Collector service with Windows Collection configured and enabled

To create an identity feed:

1. Add a destination for the feed.
 - a. Go to  (Admin) > **Services** and in the **Services**.
 - b. In the list of services, select a **Log Collector** service, and select  **View > Config**.
 - c. Select the **Event Destinations** tab.
 - d. In the Select **Event Destinations** field, select **Identity Feed**.
 - e. Click  and enter a unique name for the feed.
The Queue name identifies the feed within the Log Collector. Use the name of the feed for the Queue.
 - f. Click **OK**.
2. Test generation of messages.
 - a. Have users log into Windows boxes on the domain to generate the appropriate log messages on the domain controllers for testing.
 - b. Verify that data is written to the feed files. SSH to the Log Decoder/Collector or Virtual Log Collector being configured. Navigate to `/var/netwitness/logcollector/runtime/identity-feed` and verify that the `Identity_deploy` files are getting populated with data.
 - c. Open up a web browser (Non-Internet Explorer browsers preferred) and log in to the REST interface of the Log Collector. Use administrative credentials when logging in. For example, if the IP address of your Log Collector is 192.168.99.66, the URL would be:
 - SSL not enabled: **http://192.168.99.66:50101/event-processors**
 - SSL enabled: **https://192.168.99.66:50101/event-processors**

The browser screen should look like this:

The screen contains the name of the identity feed you created earlier (`infonetd_domain`, in this example).

For the identity feed to function correctly, port 50101 must be active on the Log Collector, and you must determine whether SSL encryption is active.

- d. Go to  (Admin) > Services > <Log Collector being setup>  > View > Explore.
- e. In the left pane, expand **rest** > **config**.


For REST to be active, **enabled** must be set to **1**.

- f. Note the value for **ssl**. If SSL should be enabled for your environment, this must be set to **on**.

Note: If you changed the setting for either the **enabled** or **ssl** option you must restart the Log Collector service before moving forward.

3. Go to  (Configure) > Custom Feeds.


The Feeds dialog is displayed.



Name	Trigger	Disk Usage	Created	Last Run Time	Status	Progress
TEST	Once	-	2019-07-13 13:30:36	2019-07-13 13:30:36	Completed	
TEST2	Once	-	2019-07-13 13:50:33	2019-07-13 13:50:33	Completed	
te	Once	-	2019-07-14 04:16:51	2019-07-14 04:16:51	Completed	
onlydom	Once	-	2019-07-14 04:21:37	2019-07-14 04:21:37	Completed	
PCAP	Once	-	2019-07-14 09:30:49	2019-07-14 09:30:49	Completed	

4. In the toolbar, click .

The Setup Feed dialog is displayed.



5. Make sure **Identity Feed** is selected and click **Next**.

The Configure Identity Feed panel opens with the **Define Feed** tab displayed.

6. (Conditional) You can create an on-demand or recurring feed.

- To define an on-demand Identity feed task that executes once, select **Adhoc** in the **Feed Task Type** field, type the feed **Name**, and browse for and open the feed.
- To define a recurring Identity Feed task that executes on a recurring basis, select **Recurring** in the **Feed Task Type** field.

The **Define Feed** dialog includes the fields for a recurring feed.

Note: RSA NetWitness Platform verifies the location where the file is stored, so that NetWitness Platform can check for the latest file automatically before each recurrence.

7. Enter a value and verify the URL field.
 - a. In the **URL** field, enter the URL where the feed data file is located. This is the REST API interface that was setup earlier. Make sure you have the following information to construct the URL:
 - The IP address of the Log Collector being used to construct the Identity Feed file.
 - The identity queue name, as set in [step 2c](#).
 - Whether or not SSL is enabled on the Log Collector REST port, as set in [step 2f](#).

You can construct this value as follows:

- SSL enabled: `https://<LogCollector>:50101/event-processors/<ID Event processor name>?msg=getFile&force-content-type=application/octet-stream&expiry=600`
- SSL not enabled: `http://<LogCollector>:50101/event-processors/<ID Event processor name>?msg=getFile&force-content-type=application/octet-stream&expiry=600`

So, using the example from earlier, the complete value that you would enter into this field is as follows:

```
http://192.168.99.66:50101/event-processors/infonetd_
domain?msg=getFile&force-content-type=application/octet-
stream&expiry=600?msg=getFile&force-content-type=application/octet-
stream&expiry=600
```

- b. For the URL verification to work correctly, it is important that the NetWitness Platform UI server can access the Log Collector's REST API port (50101). This can be tested by going to the NetWitness Platform UI server via SSH. Once there, run the following command:

- SSL enabled: `curl -vk https://<ip of log collector>:50101`
- SSL not enabled: `curl -v http://<ip of log collector>:50101`

If the `curl` command does not connect then there may be a network firewall or routing issue between the NetWitness Platform UI server and the Log Collector.

Example of a bad connection:

```
* About to connect() to 192.168.99.66 port 50105 (#0)
* Trying 192.168.99.66... No route to host
* couldn't connect to host
* Closing connection #0
curl: (7) couldn't connect to host
```

Example of a good connection:

```
* About to connect() to 192.168.99.66 port 50105 (#0)
* Trying 192.168.99.66... connected
* Connected to 192.168.99.66 (192.168.99.66) port 50105 (#0)
> GET / HTTP/1.1
> User-Agent: curl/7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7
NSS/3.19.1 Basic ECC zlib/1.2.3 libidn/1.18 libssh2/1.4.2
> Host: 192.168.99.66:50105
> Accept: */*
>
< HTTP/1.1 401 Unauthorized
< Content-Length: 71
< Connection: Keep-Alive
< Pragma: no-cache
< Expires: -1
< Cache-Control: no-cache, no-store, must-revalidate
< WWW-Authenticate: Basic realm="NetWitness"
< Content-Type: text/xml; charset=utf-8
<
<?xml version="1.0" encoding="utf-8"?>
<error>401 Unauthorized</error>
* Connection #0 to host 192.168.99.66 left intact
* Closing connection #0
```

8. The REST API requires a username and password when attempting to pull the `identity_deploy.csv` file from the Log Collector. This can be any username and password that is available on the service itself. For more information, see the "Services Security View" topic in the *Hosts and Services Guide*.

To see which accounts are available, go to  (Admin) > Services > <log collector being setup> > Actions > View > Security.

Under the Users table, you see all the users that can be used in this step. It is suggested that a separate user account is created specifically for this setup, and is used nowhere else in the environment, for added security. For details, see "Add a User and Assign a Role" in the *System Security and User Management Guide*. (Go to the [Master Table of Contents](#) to find all RSA NetWitness Platform 11.x documents.)

9. To define the recurrence interval, do one of the following:
 - Specify the number of minutes, hours, or days between recurrences of the feed.
 - Enter the date range for the execution of the feed to recur, specify the **Start Date** and time and the **End Date** and time.

10. If using SSL encryption, you need to install the REST API SSL certificate for the Log Collector into the NetWitness Platform UI server. For more information, see [Import the SSL Certificate](#).

If, after importing the SSL certificate, the verification of the URL still fails, see [Cannot Verify Identity Feed URL](#).

11. Click **Verify** to verify your identity feed configuration before you proceed to the Select Services dialog.
12. Click **Next**.

The Select Services dialog is displayed.

13. To identify services on which to deploy the feed, select one or more Decoders and Log Decoders and click **Next**.
14. Click the **Groups** tab, select a group, and click **Next**.

The Review dialog is displayed.

Note: If a group of devices with Decoders and Log Decoders is used to create recurring or custom feeds and this group is deleted, you can edit the feed and add a new group to the feed.

15. Anytime before you click **Finish**, you can:
 - Click **Cancel** to close the wizard without saving your feed definition.
 - Click **Reset** to clear the data in the wizard.
 - Click **Next** to display the next form (if not viewing the last form).
 - Click **Prev** to display the previous form (if not viewing the first form).
16. Review the feed information, and if correct, click **Finish**.

Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file are listed in the Feed grid and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.

Feeds							
☐	Name	Trigger	Feed Size	Created	Last Run Time	Status	Progress
☐	FILEHASH	Fetches STIX feeds from 2020-May-19 03:16, running every 5 minutes	0 bytes	2020-05-19 03:25:23	2020-05-22 05:16:01	Completed	
☐	FILESTIX	Fetches STIX feeds from 2020-May-19 03:32, running every 5 minutes	0 bytes	2020-05-19 03:32:09	2020-05-22 05:12:09	Completed	
☐	AllIndicatorsREST	Fetches STIX feeds from 2020-May-19 04:48, running every 5 minutes	0 bytes	2020-05-19 05:03:52	2020-05-22 05:13:26	Completed	
☐	AllIndEdited	Fetches STIX feeds from 2020-May-19 05:13, running every 5 minutes	0 bytes	2020-05-19 05:13:54	2020-05-22 05:13:54	Completed	
☐	TAXIServer1	Fetches STIX feeds from 2020-May-19 05:44, running every 5 minutes	288 bytes	2020-05-19 05:44:38	2020-05-22 05:14:38	Completed	

Import the SSL Certificate

If SSL is configured on the Identity feed's Log Collector, follow these steps to import the Log Collector's SSL certificate into the NetWitness Platform UI server key store. If this certificate is not imported, the NetWitness Platform UI server will be unable to pull the Identify feed file from the Log Collector.

1. To pull the SSL certificate off the Log Collector, SSH into the NetWitness Platform UI server and run the following command:

```
echo -n | openssl s_client -connect <HOST>:<PORT> | sed -ne '/-BEGIN
CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/<SERVERNAME>.cert
```

This command saves the SSL certificate to /tmp/<SERVERNAME>.cert. For example:

```
echo -n | openssl s_client -connect 192.168.99.66:50101 | sed -ne '/-BEGIN
CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/logcollector.cert
```

2. To import the SSL certificate into the NetWitness Platform UI server, SSH into the UI server and run the following command:

```
keytool -importcert -alias <name an alias for the cert> -file <the cert
file pathname> -keystore /etc/pki/java/cacerts
```

For example:

```
keytool -importcert -alias logcollector01 -file /tmp/logcollector.cert -
keystore /etc/pki/java/cacerts
```

3. The system requests a password. Enter the password for the keystore on the NetWitness Platform UI server, not for the jetty keystore. The default password is **changeit**.
4. Restart **jettysrv** to allow jetty to read the new certificate in the store.

Cannot Verify Identity Feed URL

If the Identity feed URL cannot be verified, and you are using SSL, make sure you followed the steps in [Import the SSL Certificate](#).

If there are issues, it is possible that the internal name of the certificate does not match the hostname of the Log Collector. The following procedure checks this.

1. SSH to the NetWitness Platform UI server.
2. Run the following command to output the CN name of the SSL cert:

```
echo -n | openssl s_client -connect <log decoder>:50101 | sed -ne '/BEGIN
CERTIFICATE-/,/-END CERTIFICATE-/p'
```

For example:

```
echo -n | openssl s_client -connect salogdecoder01:50101 | sed -ne '/BEGIN
CERTIFICATE-/,/-END CERTIFICATE-/p'
```

3. Retrieve the CN name of the SSL certificate.
4. Edit the /etc/hosts file and add the IP address and CN name to the file.
5. Restart the network service on the appliance.
6. Confirm that the name placed in the /etc/hosts file is used instead of the FQDN or IP address in the Identity feed URL.
7. Re-verify the Identity feed URL.


Edit, Upload, or Remove a Feed

You can upload a feed, edit an existing feed, or remove a feed.

To edit an existing feed:

1. Go to  (Configure) > **Custom Feeds**.


The Feeds view is displayed.



☐	Name	Trigger	Disk Usage	Created	Last Run Time	Status	Progress
☐	TEST	Once	-	2019-07-13 13:30:36	2019-07-13 13:30:36	Completed	
☐	TEST2	Once	-	2019-07-13 13:50:33	2019-07-13 13:50:33	Completed	
☐	te	Once	-	2019-07-14 04:16:51	2019-07-14 04:16:51	Completed	
☐	onlydom	Once	-	2019-07-14 04:21:37	2019-07-14 04:21:37	Completed	
☐	PCAP	Once	-	2019-07-14 09:30:49	2019-07-14 09:30:49	Completed	

2. In the toolbar, select a feed and click .

The Configure Custom Feed or Configure Identity Feed panel opens in the Custom Feed wizard.



Configure a Custom Feed

Define Feed | Select Services | Define Columns | Review

Feed Type: ☐ CSV ☒ STIX

Feed Task Type: ☒ Adhoc ☐ Recurring

Name *:

Upload As Csv File Feed: ☐

File *:

[download file](#)

— ☒ Advanced Options —


3. If you want to edit the feed file:

- a. Click **download file**.
For an Identity feed, the .zip file is downloaded. For a custom feed, the .csv or .xml file is downloaded to your local file system. For a STIX feed, the .xml file is downloaded to your local file system.
 - b. Edit and save the file.
 - c. In the **Define Feed** tab, browse for and open the edited file.
4. Edit any other parameters in the **Define Feed** tab, **Select Services** tab, and **Define Columns** tab that apply to the type of feed.
 5. Anytime before you click **Finish**, you can:
 - Click **Cancel** to close the wizard without saving your changes.
 - Click **Reset** to clear the data in the wizard.
 - Click **Next** to display the next form (if not viewing the last form).
 - Click **Prev** to display the previous form (if not viewing the first form).
 6. In the **Review** tab, review the feed information, and if correct, click **Finish**.

The feed is recreated with the updated file and new feed specifications. The feed is added to the Feeds list and progress bar tracks completion. Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file is listed in the Feeds list. You can expand or collapse the entry to see how many services are included, and which services are successful.

To upload a feed to a Decoder or Log Decoder:

1. Go to  (Admin) > **Services**.
2. Select a Decoder or Log Decoder service and click   > **View** > **Config**.
The Services Config view is displayed with the General tab open.
3. Select the **Feeds** tab.
4. In the Feeds tab toolbar, click  **Upload**.
The Upload Feeds dialog is displayed.




5. In the **File** grid, click  and select a feed file. Supported files are *.feed, *.token, and *.filter.
6. Select the feed file from the **File** list and click **Upload**.

The Upload Job list is updated to show the progress and status of the uploaded feed.

To remove a feed:


1. Go to  (Configure) > **Custom Feeds**.

The Custom Feeds view is displayed.



2. In the toolbar, select a feed and click .

The Delete Custom Feeds dialog is displayed.



You can select one of the following options to delete the feed:


- If you choose to delete the feed from **Local DB and Service**, the feed is deleted from both the service and the local NetWitness Platform box. The deleted feed will no longer be seen on the NetWitness Platform user interface.
- If you choose to delete the feed from **Local DB Only**, the feed is deleted from the local NetWitness Platform box. The deleted feed will not be seen on the NetWitness Platform user interface; however, the last deployed version of the feeds will be present on the service. The undeployed feeds will be deleted forever.
- If you choose to delete the feed from **Service Only**, the feed is deleted from the service. The deleted feed will appear on the NetWitness Platform user interface and can be deployed again.

3. Select where you want to delete the feed and click **Delete**.

A warning dialog is displayed.

4. Click **yes** to confirm that you want to delete the feed from the select areas.

- If you chose to delete the feed from the **Local DB Only**, the feed is deleted.
- If you chose to delete the feed from the **Local DB and Service** or **Service Only**, the Delete Custom Feeds view is displayed showing the progress of the deletion on the service.



Create Custom Meta Keys Using a Custom Feed

This topic provides information on how to add custom meta keys using a custom feed in the Log Decoder, and highlights the configuration changes to reflect the custom meta keys in the Concentrator, ESA, Archiver, Warehouse Connector, and Reporting Engine schema. You can create custom meta keys to retrieve data, to investigate and analyze the logs and packets. Custom meta keys enable you to add an enrichment context for the log and packet data.

Here is an example of creating a custom meta key in the Log Decoder. In this scenario, an organization wants to track the location of an asset such as a printer. So, a custom meta key **source location** is introduced, which indicates the location of the asset, for example Printer1, which is located in the 'Fifth Floor A wing'.

Note: Custom meta keys can be created in the Decoder as well. Select the `index-decoder-custom.xml` file when you create a custom meta key in the Decoder.

Add a Custom Meta Key in the Log Decoder

To add custom meta keys using custom feed:

1. Go to  (Admin) > **Services**.
2. Select a Log Decoder service and click   > **View** > **Config** > **Files** tab > **index-logdecoder-custom.xml**.

```
<Language>
  <?xml version="1.0" encoding="utf-8"?>
  <Language level="IndexNone" defaultAction="Auto">
    <!-- Reserved Meta key for Feed -->
    <Key description="Source Location" level="IndexNone"
name="location.src" format="Text"/>
  </Language>
  name = Name of the key (max is 16 chars)
```

3. Restart the Log Decoder service. In the Services view, click   > **Restart**.

Note: Following are the NetWitness Platform reserved keywords. Do not use these keywords as custom meta keys as it can cause display error in the NetWitness Platform user interface due to CSS conflicts referencing to the same keywords.

risky
safe
unsafe
ecat
black
ecat-text
bold
machine-name
critical
item
row
list
ioc-info
status-info
details

```
score
cards
intro-subject
context-intro
im-risk-score-type
im-priority-type
im-score-type
incident-details
cs-name
im-score-critical
im-score-high
im-score-medium
im-score-low
im-score-lowest
im-score-neutral
alert-score-high
ecat-score
loading-msg
content-loading
live-title
live-text
```

Deploy a Log Decoder Feed in Live

To deploy the feed in the live environment:

1. Go to  (Configure) > **Live Content**
2. Select a group of resources, or a previously created resource package. To select a resource or group of resources:
 - a. In the **Live Search View**, browse Live resources (for example, search for the **Log Collector** resource Type).
 - b. In the **Matching Resources** panel, select **Show Results > Grid**.

- c. Select the checkbox to the left of the resources that you want to deploy.

- d. In the Matching Resources toolbar, click Deploy.
3. To select a resource package to deploy:
- a. In the **Live Search** view - **Matching Resources** toolbar, select **Package > Deploy** :
The Package page of the Resource Package Deployment wizard is displayed.

- b. Click **Browse** and select a package from your network.
- c. Click **Open**.
At this point, whether you are deploying a package or a group of resources, the Deployment Wizard opens, and the Resources page is displayed.

3. Click **Next**.


The **Services** page is displayed. It has two tabs, **Services** and **Groups**, which provide a list of services and service groups that are configured in the  (Admin) > Services view. The columns are a subset of the columns available in the Services view.

Note: The Live server is "smart" about deploying resources to Services. For example, it does not deploy resources that have a Medium of packets to any Log Decoders. This means that only applicable content resources are deployed to each Service.

4. Select the services to which you want to deploy the content. You can select any combination of services and service groups.

Use the **Services** tab to select individual services, list of services and service groups that are configured in the Admin Services view.

Use the **Groups** tab to select groups of services.




The screenshot shows the 'Deployment Wizard' interface with the 'Services' tab selected. The wizard has four steps: Resources, Services (current), Review, and Deploy. Below the tabs, there are two sub-tabs: 'Services' (selected) and 'Groups'. A table lists available services with checkboxes for selection. The table has columns for 'Name' and 'Type'.

	Name ^	Type
☐	UI Endpoint	Other

At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Next'.

5. Click **Next**.

The **Review** page is displayed.



The screenshot shows the 'Deployment Wizard' interface with the 'Review' tab selected. The wizard has four steps: Resources, Services, Review (current), and Deploy. Below the tabs, there is a table summarizing the selected services and resources.


Service	Service Type	Resource Name	Resource Type
UI Endpoint	NW Local	Basic Rule Template	RSA Event Stream Analy...

At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Deploy'.


Make sure that you have selected correct resources and the services to which you want to deploy them.

6. Click **Deploy**.

The **Deploy** page is displayed. The Progress bar turns green when you have successfully deployed the resources to the selected services.



If you try to deploy resources and services that are not compatible, NetWitness Platform displays the Errors and Retry buttons, which you click to review the errors and retry the deployment.



7. Click **Close**.

Note: The Source IP should be indexed by selecting the type as 'IP' as the ip.src. and ip.dst are in IPv4 format.

In this scenario, a custom meta key location.src (location source) is added by indexing the hostname (alias.host). In this example, the printer hostname are populated in meta key 'alias.host'. Select **alias.host** as callback key, and index type as 'Non IP' in the Feed Wizard as shown below. In the Define Values section, select the custom meta key from the drop down menu.

Add the Custom Meta Key Entry in the Concentrator Custom Index file

To add the custom meta key entry in the Concentrator custom index file:

1. Go to  (Admin) > **Services** > **Concentrator**.
2. Click  > **View** > **Config** > **Files** tab > **index-concentrator-custom.xml**.
3. Add the custom meta key entry in the Concentrator index file.

```
<Language>
  <?xml version="1.0" encoding="utf-8"?>
  <Language level="IndexNone" defaultAction="Auto">
    <!-- Reserved Meta key for Feed -->
    <Key description="Source Location" level="IndexValues"
      name="location.src" format="Text" valueMax="10000"
      defaultAction="Open"/>
  </Language>
```

4. To restart the Concentrator service, in the Services view, click  > **Restart**.


Note: In case of the Broker, the Broker derives its index from the Concentrator from which it aggregates. So you do not need to create custom meta in the Broker. If you have not indexed the meta key in the Concentrator, the Broker does not display the meta key in Investigate.

Investigate a Custom Meta Key

Note: You have to log out and log back in to the NetWitness Platform User Interface in order to view the custom meta key in Investigate.

To investigate a custom meta key:

1. Go to **Investigate**.
The Investigate dialog, which provides services, is displayed.
2. Select a Concentrator service, and click **Navigate**.




Additional Procedures

The following procedures must be executed if you have Warehouse Connector, Archiver, Reporting Engine, and ESA configured.

Verify the Custom Meta Keys on ESA

After you add custom meta keys on the Concentrator, you can verify that your meta keys are updated on ESA.

1. Go to  (Configure) > **ESA RULES** > **Settings** tab.



2. In the Meta Key References, click the Meta Re-Sync (Refresh) icon (.
3. Verify that the custom meta keys appear on ESA. If you do not see the meta keys, you may need to restart the Concentrator.

Update the Schema in Archiver

If you want to configure the Archiver, using the new custom meta keys, you need to update the Archiver schema in the Reporting Engine. To update the Archiver schema in Reporting Engine:

1. Go to  (Admin) > **Services** > **Archiver**.
2. Select  > **View** > **Config** > **Files** > **index-archiver-custom.xml**.
3. Add the custom meta key entry in the Archiver index file.

```
<Language>
  <?xml version="1.0" encoding="utf-8"?>
  <Language level="IndexNone" defaultAction="Auto">
    <!-- Reserved Meta key for Feed -->
    <Key description="Source Location" level="IndexValues"
      name="location.src" format="Text"
      valueMax="10000" defaultAction="Open"/>
  </Language>
```

4. To restart the Archiver service, click  > **Restart**.
The Archiver schema is updated with the custom meta key.

Update the Schema in Warehouse Connector

If you want to configure the Warehouse Connector with custom metadata and use it in a Warehouse Connector report then you need to update the Warehouse Connector schema in the Reporting Engine.


If the Log Decoder or Decoder, where the custom meta key is added, is one of the sources in the Warehouse Connector stream, you need to update the schema in the Warehouse Connector.

To update the Warehouse Connector schema in the Reporting Engine:

1. Go to  (Admin) > **Services** > **Warehouse Connector**.
2. Click  > **View** > **Config**.
The Services Config view of Warehouse Connector is displayed.
3. Click the **Streams** tab.

4. Select the stream and then click **Reload**.

The Warehouse Connector pulls the schema from the downstream devices (Log Decoder/Decoder).



For more information on streams, see "Configure Streams" in the *Warehouse Connector Configuration Guide*.

Update the Schema in Reporting Engine

To update the schema in Reporting Engine:

1. Go to  (Admin) > **Services** > **Reporting Engine**.
2. Click  > **Restart**.

Note: Restart the Reporting Engine or wait for thirty minutes for the schema to be updated.

To view the custom meta key:

1. Navigate to **Monitor** > **Reports** > **Rules**.
2. In the toolbar, click .
3. Select **Warehouse DB**.
4. In the **Build Rule** tab, search for the custom meta from the right panel.
The custom meta key is displayed.

ManageView[RULE] New Rule

Build Rule

Rule TypeWarehouse DB

Expert Mode☐

NameCustom Meta

Selectloc

Fromloc_cityloc_countryloc_descloc_statelocation_srclog_session_idlog_session_id1logon_type7.71351e+31longdec_dst4.86134e+30longdec_src

WhereSource Location

Group By

Having

Order By

Limit20

Meta

Warehouse

locat

location_src

Lists

Filter

Insert

☐ Attack Kill Chain Report☐ Compliance☒ Critical Windows Machines☐ Dev☒ Infected Filenames from ECAT☒ Local_Country

Decoder and Log Decoder Additional Procedures

This topic explains the additional procedures an administrator could choose to follow which are not essential for the configuration of the Decoder or Log Decoder.

Topics

- [Configure High Speed Packet Capture Capability \(Version 11.6 and Later\)](#)
- [Configure a Log Decoder to Accept Protobuf](#)
- [Configure Session Split Timeouts](#)
- [Configure Syslog Forwarding to Destination](#)
- [Configure Transaction Handling on a Decoder](#)
- [Configure Data Export](#)
- [Decrypt Incoming Packets](#)
- [Edit Decoder System Configuration](#)
- [Enable CPU Usage Statistics for Installed Content](#)
- [Enable Parser Mappings](#)
- [Enable or Disable Lua and Flex Parsing Systems](#)
- [Map IP Address to Service Type for Log Parsing](#)
- [Map an IP Address to a Time Zone](#)
- [Obtain Log Files from a Pre-11.0 Log Decoder](#)
- [Upload a Log File to a Log Decoder](#)
- [Upload a Packet Capture File](#)

Configure 10G Capability (Version 11.5 and Earlier)

This topic guides administrators in how to tune a Network Decoder specifically for high speed packet capture using NetWitness Platform 11.x. This applies when capturing packets on a 10G interface card.

Caution: Packet capture at high speeds requires careful configuration and pushes the Decoder hardware to its limits, so please read this entire topic when implementing a 10G capture solution.

Note: Enabling intra-session for HTTP pipelining is not supported when capturing at 10G rates on a single Decoder, as the HTTP_Lua parser required to function can cause dropped packets at that rate of ingest. In this case, the load should be spread across multiple Concentrator-Decoder pairs.

RSA NetWitness Platform provides support for high-speed collection on the Decoder. You can capture network packet data from higher speed networks and optimize your Network Decoder to capture network traffic up to 8Gb/sec sustained and 10Gb/sec burst, depending on which parsers and feeds you have enabled.

Enhancements that facilitate capture in these environments include the following:

- Utilization of the `pf_ring` capture driver capability to leverage the commodity 10G Intel NIC card for high-speed capture.
- Introduction of `assembler.parse.valve` configuration, which automatically disables application parsers when certain thresholds are exceeded, to limit risk of packet loss. When the application parsers are disabled, network layer parsers are still active. When stats fall below exceeded thresholds, application parsers are automatically re-enabled.
- Utilization of Berkeley Packet Filters (BPF) in 10G environments. For information about using BPF, see "(Optional) Configure System-Level (BPF) Packet Filtering" in [Configure Capture Settings](#).

Note: Beginning with 11.5, the Network Decoder can capture from multiple interfaces simultaneously. This functionality allows Network Decoders to capture from multiple physical Network Interface Cards (NICs) while leveraging the same network rules, application rules, and parsers for each NIC. You can use this feature in 10G environments. For more information, see [\(Optional\) Multiple Adapter Packet Capture](#).

Hardware Prerequisites

- A Series 4S, Series 5 or Series 6 Decoder
- An Intel 82599-based ethernet card, such as the Intel x520. All RSA-provided 10G cards meet this requirement. Two examples are:
 - All SMC-10GE cards provided by RSA.
 - A Dell Network Daughter Card using an Intel controller to provide 10G network interfaces. This is included in all Series 5 hardware.
- For the Series 4S / Dell R620 only: 96 GB of DD3-1600 memory in **dual-rank** DIMMs. Single-rank DIMMs may decrease performance by as much as 10%. To determine the speed and rank of the installed DIMMs, run this command:

```
dmidecode -t 17.
```
- Sufficiently large and fast storage to meet the capture requirement. Storage considerations are covered later in this topic.
- Each Network Decoder configured with a minimum of 2 DACs or SAN connectivity.

Software Prerequisites

- Dell R620-based systems, such as the Series 4S, must have their BIOS updated to v1.2.6 or later.
- The 10G Decoder capability is only supported on RSA-provided Decoder Installation images. All required software is installed by default.
- If upgrading from a previous release, perform the upgrade first before proceeding with configuration

Install the 10G Decoder

Note: You can skip to "Configure the 10G Decoder" if you are starting with new Series 5 or Series 6 hardware.

Perform the following steps to install the NetWitness 10G Decoder:

Download and Update the BIOS

Note: BIOS revisions earlier than v1.2.6 have issues properly identifying the location of the 10G capture card within the system. It is recommended that customers update to the latest v2.2.3 BIOS, but is not required for 10G if they are running v1.2.6 or later.

1. Download BIOS v2.2.3 from the following location:
<http://www.dell.com/support/home/us/en/19/Drivers/DriversDetails?driverId=V7P04>
2. Download the Update Package for the Red Hat Linux file.
3. Copy the file to the NetWitness server.
4. Login as `root`.
5. Change the permissions on the file to execute.
6. Run the following file:

```
./BIOS_V7P04_LN_2.2.3.BIN
```
7. Reboot the system when execution is complete and a reboot is requested.

Note: The BIOS installation procedure takes approximately 10 minutes.

Locate the 10G Decoder Packages

The packages required to configure the 10G Decoder should already be present on the Decoder installation image. You should not have to install any additional packages.

Verify 10G Decoder Packages Are Installed

Installation of the 10G Decoder packages is handled automatically. Therefore, there should be no action to enable the 10G functionality.

- If you upgraded the kernel packages as part of an upgrade, a reboot is required. The operating system will recompile and install the drivers for the upgraded kernel.
- You can verify that the installation was successful if you see additional `PFRINGZC` interfaces available when selecting the Capture Port Adapter as described below.

Configure the 10G Decoder

Perform the following steps to configure the 10G Decoder:

1. From the **Decoder Explore** view, right-click **Decoder** and select **Properties**.
 2. In the properties drop-down menu, select **reconfig** and enter the following parameters:
`update=1 op=10g`
 This adjusts the Decoder packet processing pipeline to allow for higher raw data throughput, but less parsing ability.
 3. From the **Decoder Explore** view, right-click **database** and select **Properties**.
 4. In the **Properties** drop-down menu, select **reconfig** and enter the following parameters:
`update=1 op=10g`
 These parameters adjust the packet database to use very large file sizes and Direct I/O.
 5. To capture traffic on a single port,
 - a. Go to **(missing or bad snippet)>Services**, select the target Network Decoder type and click **>View>Config**.
 - b. In Decoder Configuration section, in **Capture Interface Selected** field, select the interface. For example, `PFRINGZC,p1p1` or `PFRINGZC,p1p2`.
 6. To capture traffic on two ports simultaneously on the same 10G card,
 - a. Go to **(missing or bad snippet)>Services**, select the target Network Decoder type and click **>View>Config**.
 - b. In the Decoder Configuration section, select the first interface in **Capture Interface Selected** field. For example, `PFRINGZC,p1p1` or `PFRINGZC,p1p2`.
 - c. Select **Explore** from the drop-down menu and click `decoder> config` node. In the **capture.device.params** field, enter the second interface variable that used on this 10G card for simultaneous capture. For example, `device=zc:p1p2,zc:p1p1`
 7. If the write thread is having trouble sustaining the capture speed, you can try the following:
 Change `/datebase/config/packet.integrity.flush` to `normal`.
- Note:** You can adjust the `packet.file.size` to a higher value, but keep the file size under 10 GB, as the whole file is buffered in memory.
8. (Optional) Application parsing is extremely CPU intensive and can cause the Decoder to drop packets. To mitigate application parsing-induced drops, you can set `/decoder/config/assembly.parse.valve` to `true`. These are the results:
 - When session parsing becomes a bottleneck, application parsers (HTTP, SMTP, FTP, and others) are temporarily disabled.
 - Sessions are not dropped when the application parsers are disabled, just the fidelity of the parsing performed on those sessions.
 - Sessions parsed when the application parsers are disabled still have associated network meta (from the network parser).
 - The statistic `/decoder/parsers/stats/blowoff.count` displays the count of all sessions that bypassed application parsers (network parsing is still performed).

- When session parsing is no longer a potential bottleneck, the application parsers are automatically re-enabled.
 - The assembler session pool should be large enough that it is not forcing sessions.
 - You can determine if sessions are being forced by the statistic `/decoder/stats/assembler.sessions.forced` (it will be increasing). Also `/decoder/stats/assembler.sessions` will be within several hundred of `/decoder/config/assembler.session.pool`.
9. (Optional) If you need to adjust the MTU for capture, add the `snaplen` parameter to `capture.device.params`. Unlike previous releases, the `snaplen` does not need to be rounded up to any specific boundary. The Decoder automatically adjusts the MTU set on the capture interfaces.
 10. The following configuration parameters are deprecated and no longer necessary
 - The `core=` parameter in `capture.device.params`
 - Any configuration files under `/etc/pf_ring` directory

Note: An Ethernet device installed post imaging does not require any configuration for use as a capture device. It does require configuration if it is used as a network interface, or for system tools to access it without manual configuration.

Typical Configuration Parameters

Typical configuration parameters are listed below. Actual parameters may vary depending on the amount of memory and CPU resources available.

1. Session and packet pool settings (under `/decoder/config`):
 - `pool.packet.pages = 1000000`
 - `pool.session.pages = 300000`
2. Packet write block size under (`/database/config/packet.write.block size`) set to `filesize`.

Note: This configures the Decoder to buffer the file with huge pages and write using direct I/O for maximum performance.

3. Parse Thread Count (under `/decoder/config`).


```
parse.threads =12
```

Storage Considerations

When capturing at 10G line rates, the storage system holding the packet and meta databases must be capable of sustained write throughput of 1400 MBytes/s.

Using the Series 4S Hardware (With Two or More DAC Units)

The Series 4S is equipped with a hardware RAID SAS controller capable of an aggregate 48Gbit/s of I/O throughput. It is equipped with eight external 6 Gbit ports, organized into two 4-lane SAS cables. The recommended configuration for 10G is to balance at least two DAC units across these two external connectors. For example, connect one DAC to one port on SAS card, and then connect another DAC to the other port on the SAS card.

For environments with more than two DACs, chain them off each port in a balanced manner. This may require re-cabling of DACs in an existing deployment, but should not affect data that has already been captured on the Decoder.

Note: If you are adding new storage capacity, refer to "Configure Storage Using the REST API" in the *Storage Guide for RSA NetWitness Platform*. (missing or bad snippet)

Using SAN and Other Storage Configurations

The Decoder allows any storage configuration that can meet the sustained throughput requirement. The standard 8-Gbit FC link to a SAN is not sufficient to store packet data at 10G; in order to use a SAN it may be required to perform aggregation across multiple targets using a software-RAID Scheme. Thus environments using SAN are required to configure connectivity to the SAN using multiple FCs.

Capturing 10G with 40G Network Fabric

The Network Decoder is capable of physically connecting to a 40Gbps network infrastructure if you use the applicable hardware described in the *Series 6 Hardware Setup Guide*, located here: <https://community.rsa.com/community/products/netwitness/hardware-setup-guides>.

To support this feature without causing significant packet drops, the following conditions must be met:

- The capture rate must be throttled to 10Gbps by the network.
- The guidelines described later in this topic in [Tested Live Content](#) still apply.

Parsing and Content Considerations

Parsing raw packets at high speeds presents unique challenges. Given the high session and packet rates, parsing efficiency is paramount. A single parser that is inefficient (spends too long examining packets) can slow the whole system down to the point where packets are dropped at the card.

For initial 10G testing, start with only native parsers (except SMB/WebMail). Use the native parsers to establish baseline performance and with little to no packet drops. Do not download any Live content until this has been done and the system is proven to capture without issue at high speeds.

After the system has been operational and running smoothly, Live content should be added very slowly - especially parsers.

Best Practices

Whether you are updating a currently deployed system or deploying a new system, it is recommended you use the following best practices to minimize risk for packet loss. One caveat is if you are updating a current 10G deployment but not adding any additional traffic. For example, a current Decoder capturing off a 10G card at 2G sustained should see no difference in performance, unless part of the update also entails adding additional traffic for capture.

- Incorporate baseline parsers (except SMB/Webmail, both of which generally have high CPU utilization) and monitor to ensure little to no packet loss.
- When adding additional parsers, add only one or two parsers at a time.
- Measure performance impact of newly added content, especially during peak traffic periods.
- If drops start occurring when they did not happen before, disable all newly-added parsers and enable just one at a time and measure the impact. This helps pinpoint individual parsers causing detrimental effects on performance. It may be possible to refactor it to perform better or reduce its feature set to just what is necessary for the customer use case.
- Although lesser performance impacts, feeds should also be reviewed and added in a phased approach to help measure performance impacts.
- Application Rules also tend to have little observable impact, though again, it is best not to add a large number of rules at once without measuring the performance impact.
- If you continually get a timeout message in the Investigate > Events view, such as `The query on channel 192577 was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS)`, first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or other warnings that may need to be addressed to increase query response time. If there are no messages indicating any specific warnings then try increasing the Core Query Timeout from the default 5 minutes to 10 minutes as described in "View Query and Session Attributes per Role" section of the *System Security and User Management Guide*.

Finally, making the recommended configuration changes outlined in the Configuration section will help minimize potential issues.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- MA content (7 Lua parsers, 1 feed, 1 application rule)
- 4 feeds (alert ids info, nrmalwaredomains, warning, and suspicious)
- 41 application rules
- DNS_verbose_lua (disable DNS)
- fingerprint_javascript_lua
- fingerprint_pdf_lua
- fingerprint_rar_lua
- fingerprint_rtf_lua
- MAIL_lua (disable MAIL)
- SNMP_lua (disable SNMP)
- spectrum_lua
- SSH_lua (disable SSH)

- TLS_lua
- windows_command_shell
- windows_executable

NOT TESTED:

- SMB_lua, native SMB disabled by default
- html_threat

OTHER:

- HTTP_lua reduces the capture rate from >9G to <7G. At just under 5G this parser can be used in place of the native without dropping (in addition to the list above).
- xor_executable pushes parse CPU to 100% and the system can drop significantly due to parse backup.

Aggregation Adjustments Based on Tested Live Content

A 10G Decoder can serve aggregation to a single Concentrator while running at 10G speeds. Deployments using Malware Analysis, Event Stream Analysis, Warehouse Connector, and Reporting Engine are expected to impact performance and can lead to packet loss.

For the tested scenario, the Concentrator aggregates between 45 and 70k sessions/sec. The 10G Decoder captures between 40-and 50k sessions/sec. With the content identified above, this is about 1.5 to 2 million meta/sec. Due to the high volume of session rates, the following configuration changes are recommended:

- Nice aggregation on the Concentrator limits the performance impact on the 10G Decoder. The following command turns on nice aggregation.
`/concentrator/config/aggregate.nice = true`
- Due to the high volume of sessions on the Concentrator, you may consider activating parallel values mode on the Concentrator by setting `/sdk/config/parallel.values` to 16. This improves Investigation performance when the number of sessions per second is greater than 30,000.
- If multiple aggregation streams are necessary, aggregating from the Concentrator instead has less impact on the Decoder.
- Further review for content and parsing is required for deployments where you want to use other NetWitness Platform components (Warehouse Connector, Malware Analysis, ESA, and Reporting Engine).

Optimize Read/Write Operations When Adding New Storage

A 10G Decoder is optimized to stagger read and write operations across multiple volumes so that the current file being written is on a different volume from the next file that will be written. This allows maximum throughput on the raid volume when reading data from the last file being written while writing the current file on a different volume. However, if volumes are added after a Decoder has been in use, the ability to stagger is limited because one or more volumes are already full so the new volume is the only place new files can be written.

To remedy this situation, an administrator can run a `stagger` command on an existing NetWitness Platform database (packet, log, meta, or session), that has at least two volumes, to stagger the files across all volumes in the most optimal read/write pattern. The major use case is when new storage is added to an existing Decoder and you want to stagger the volumes BEFORE restarting capture.

The configuration nodes for this command are the session, meta, and packet databases. Each of these lives under `/database/config`, which is usually a root node. The config nodes for a Decoder are:

- `/database/config/packet.dir`
- `/database/config/meta.dir`
- `/database/config/session.dir`

The *NetWitness Platform Core Database Tuning Guide* has information on how those configurations are formatted.

The `stagger` command is typically only useful for a 10G Decoder and usually just for the packet database. Maximum performance is achieved for storing and retrieving packets when multiple volumes are present. In this scenario, the Decoder always fills the volume with the most free space. When the volumes are roughly the same size, this results in a staggered write pattern, which allows maximum throughput for reading and writing across all volumes. However, this only naturally occurs when multiple packet storage volumes are present at the time the Decoder is first deployed.

A typical use case is adding more storage to an existing Decoder to increase retention. However, when adding storage to a deployment that has already filled the existing volumes with stored packets, the Decoder will naturally fill the new storage with packets before rolling out any packets on the existing storage. This results in a suboptimal read/write pattern because most reads will occur on the same volume that is currently being written to. In a 10G deployment, reads are blocked from the volume when writes are occurring. This does not stop ALL reads on that volume, because the file is buffered in memory before being written, but it does result in suboptimal read performance.

With the `stagger` command, you can add more storage and then have the service naturally stagger the files across ALL volumes (existing and new) so that read performance is optimized.

Caution: This command should only be performed AFTER the storage is mounted and the Decoder configured to use it (for example, after adding the mount point(s) to `packet.dir`).

The downside to this command is it can take some time to stagger and the Decoder should not be capturing during the stagger operation.

Recommended workflow:

1. Add all storage and configure mount points.
2. Add new storage mount points to `packet.dir` (or `session.dir/meta.dir`) and restart service (very important!).
3. Ensure capture is stopped.
4. Run stagger operation but make sure the connection that initiated the stagger operation is never terminated until the operation is complete. If the connection is terminated, then the stagger operation will be canceled. If the operation is canceled, the files that were already staggered will remain in place. The operation can be resumed by rerunning the same command (the work already done will not need to be done again). If running stagger from NwConsole, run the `timeout 0` command before sending the `stagger` command. This will prevent the normal 30-second command timeout.
5. Start capture after `stagger` command finishes.

The following are the parameters for the command:

- `type` - The database that will be staggered (session, meta, or packet). Typically only the packet database is useful for staggering, but it is possible to do the session or meta database when multiple volumes are present for those databases. Since the session and meta databases write far less data than the packet database, typically staggering those databases results in less noticeable performance gains.
- `dryRun` - If `true` (the default), will only return a description of the operations that would be performed. If `false`, then the files will actually be moved to an optimal read/write pattern. You **MUST** pass `false` to actually stagger the files.

Example usage from NwConsole:

```
login <decoder>:50004 <username> <password>
timeout 0
send /database stagger type=packet dryRun=false
```

If you run this command via the RESTful API, please pass the additional parameter `expiry=0` to prevent a timeout from the service. You will also need to ensure the HTTP client does not disconnect before the operation completes.

Configure High Speed Packet Capture Capability (Version 11.6 and Later)

This topic guides administrators on how to tune a Network Decoder specifically for high speed packet capture using NetWitness Platform 11.6. This applies when capturing packets on 10G, 40G, or higher speed interface cards.

IMPORTANT: Packet capture at high speeds requires careful configuration and pushes the Decoder hardware to its limits, so you must read this entire topic while implementing a high speed capture solution.

RSA NetWitness Platform provides support for high-speed collection on the Decoder. You can capture network packet data from higher speed networks and optimize your Network Decoder to capture network traffic up to 40 Gbit/s. These network capture speeds can be easily achieved on a Decoder but make a note that not all NetWitness Platform features operate at these rates. The actual throughput of Decoder depends on the following factors:

- **Amount of packets filtered vs. the amount of packets retained**

It is important to filter the traffic that enters the Decoder to minimize packet drops. You can use Network Rules or a BPF Rule to filter the traffic before it enters advanced features such as stream reassembly and deep packet inspection. For more information, see [Configure Network Rules](#) and [\(Optional\) Configure System-Level \(BPF\) Packet Filtering](#).

- **Average packet size**

Decoder capturing small packets, for example 64 byte packets, can decrease the overall throughput of the Decoder. If average packet sizes are small, the line rate of Decoder ingest will fall significantly below the network link speed. It occurs due to the amount of per-packet overhead between each Ethernet frame.

- **Shared system resources**

If the Decoder host is also hosting other applications or add-on features that manipulate data and perform other analysis, it can lower the overall throughput of the Decoder. Even if these add-on applications are not utilizing the CPU, these will utilize the shared resources like memory bandwidth and CPU interconnect bandwidth.

Enhancements that facilitate capture in these high speed environments include the following:

- Utilization of the DPDK capture drivers to leverage the commodity 10G and 40G Intel NIC cards for high-speed capture environments. For more information on how to configure DPDK, see [\(Optional\) Data Plane Development Kit Packet Capture](#).
- Introduction of `assembler.parse.valve` configuration, which automatically disables application parsers when certain thresholds are exceeded, to limit risk of packet loss. When the application parsers are disabled, network layer parsers are still active. When stats fall below exceeded thresholds, application parsers are automatically re-enabled.

- Utilization of Berkeley Packet Filters (BPF) in 10G environments. For more information, see [\(Optional\) Configure System-Level \(BPF\) Packet Filtering](#) in [Configure Capture Settings](#).
- Decoder will utilize the symmetric RSS to distribute capture loads among multiple CPU cores. For more information see, "*Utilizing Receive Side Scaling with DPDK*" in [\(Optional\) Data Plane Development Kit Packet Capture](#)
- Decoder will distribute the session reassembly work among multiple CPU cores as well, provided that the packets fed into each assembly are naturally segregated using the RSS features or as a result of being captured on different physical interfaces.

Note: Beginning with 11.5, the Network Decoder can capture from multiple interfaces simultaneously. This functionality allows Network Decoders to capture from multiple physical Network Interface Cards (NICs) while leveraging the same network rules, application rules, and parsers for each NIC. You can use this feature in all capture environments. For more information, see [\(Optional\) Multiple Adapter Packet Capture](#).

Note: Enabling intra-session for HTTP pipelining is not supported when capturing at 10G rates on a single Decoder as the HTTP_Lua parser required to function can cause dropped packets at that rate of ingest. In this case, the load should be spread across multiple Concentrator-Decoder pairs.

Hardware Prerequisites

- A Series 5 or Series 6 Decoder
- An Intel 82599-based ethernet card, such as the Intel x520 or an Intel i40e-based card like the Intel XL710. All RSA-provided 10G cards meet this requirement such as the following:
 - All SMC-10GE cards provided by RSA.
 - A Dell Network Daughter Card using an Intel controller to provide 10G network interfaces. This is included in all Series 5 hardware.
- Sufficiently large and fast storage to meet the capture requirement of packets in case you set up in a mode other than meta-only. Storage considerations are covered later in this topic.
- If the Decoder will reassemble and handle high numbers of sessions, for example 30000 sessions per second or more, then at least two physical storage volumes are required. This will allow one volume to be used for writing session data while the other to read session data and deliver it for down-stream analytics.
- Each Network Decoder configured with a minimum of 2 DACs or SAN connectivity when set up in mode other than meta-only.

Software Prerequisites

RSA NetWitness Platform 11.6 or later.

Decoder Installation

Perform the following steps to install the Decoder.

Install the Decoder Service on Host that Will Perform the Capture

The Network Decoder service package contains the necessary software components to utilize the features described in the document. The Network Decoder service package and all other required packages are installed during the Decoder installation.

Select the Capture Interfaces and Assign Interfaces to DPDK

You must choose the physical network connections that will be used for packet ingest. The physical interface type will depend on the type of network traffic sent to the Decoder. The interface you select must match the physical characteristics of the Decoder feed (or network link) used to receive packets. For example, if you want to capture at 40G speeds on a cable, then you need a 40G interface capable of attaching to that cable.

For information on how to assign interfaces to DPDK, see [\(Optional\) Data Plane Development Kit Packet Capture](#).

Verify the Decoder Packages Installation

Perform the following steps to verify the Decoder package installation:

- After you configure the interfaces to use with DPDK, ensure that you reboot the Decoder host.
- After the reboot, ensure that the DPDK interfaces appear in the Decoder's capture adapter list.

Select the Decoder Operating Mode

You can configure Decoder to handle any type of capture scenario that ranges from a very deep application-level inspection to very fast and simple network connection tracking. The Decoder includes the following three built-in templates that serve as a starting point for configuring your capture needs:

1. **Normal:** It is the default Decoder mode with no defined resource allocation and all tunable parameters set to default. It can be a good option if your capture rate is less than 5 Gbit/s or if you want to run a large amount of deep inspection content such as application parsers, feeds, OpenAppID detectors, complex Snort rules, and so on. For more information, see [Configure the Decoder for Normal Mode \(Default Mode\)](#).
2. **10G:** This mode is available in NetWitness Platform 10.4 and later. You can use this mode if your capture rate is under 10 Gbit/s. This mode assumes that most of the packets are retained and saved in the packet database. Make a note that not all parsers can run at 10 Gbit/s speeds as some parsers are too complex for high speeds. For more information, see [Configure the Decoder for 10G Mode](#).
3. **NDR:** This mode is to capture beyond the 10 Gbit/s speeds, but below 40Gb/sec. Make a note of the following points while using this mode:
 - a. By default, this mode assumes that packets are not retained as the rate at which this mode consumes storage is significantly high.
 - b. By default, the ingested packets are dropped. You must insert the network rules before the drop statement to pick portions of the network stream and keep the incoming packet feed for further analysis.

- c. Make a note that at higher speeds only a limited amount of deep packet inspection content can be used.
- d. Ensure that you enable the multi-thread assembly while using this mode.

For more information, see [Configure the Decoder for NDR Mode](#).

Configure the Decoder for Normal Mode (Default Mode)

By default, the Decoder runs in the normal mode and requires no additional configuration. The default normal mode captures up to 5Gb/sec with large amounts of deep packet inspection while storing network sessions.

Configure the Decoder for 10G Mode

The instructions in this section are to configure the Decoder for capture speeds up to 10Gb/sec with medium amounts of deep packet inspection while storing network sessions.

To configure the Decoder:

1. From the **Decoder Explore** view, right-click **Decoder** and select **Properties**. In the properties drop-down menu, select **reconfig** and enter the following parameters:
`op=10g`
 Ensure that the correct parameter is displayed in the output.
2. Add the `update=true` parameter and run **reconfig** again to save the configuration changes. For example, your final parameters can be `op=10g update=true`.
3. From the **Decoder Explore** view, right-click **database** and select **Properties**.
4. In the **Properties** drop-down menu, select **reconfig** and enter the following parameters:
`update=1 op=10g`
 These parameters adjust the packet database to use very large file sizes and Direct I/O.
5. Perform the following steps to select capture interfaces:
 - a. In the **Decoder Explore** view, right-click on **Decoder** and select **Properties**.
 - b. In the **Properties** drop-down menu, click **select** and run the command to view the available adapters. For example, this might produce output like the list below. The actual list of adapters will vary depending on the hardware available on the Decoder host.


```
1: packet_mmap_,eth0
2: DPDK,0000:41:00.0
3: DPDK,0000:41:00.1
```
 - c. To select a single interface for capture, enter the parameter `adapter=N` where N is the capture interface. For example, using the list above you might choose to capture on the first DPDK interface using the `adapter=2` parameter.
 - d. To select more than one interface for capture, enter the parameter `adapter=N,M,...` where N,M and so on are the capture interfaces. For example, using the list above you might choose to capture on both DPDK interfaces using the `adapter=2,3` parameter.

6. (Optional) Application parsing is extremely CPU intensive and can cause the Decoder to drop packets. To mitigate application parsing-induced drops, you can set `/decoder/config/asmblr.parse.valve` to `true`. These are the results:
 - When session parsing becomes a bottleneck, application parsers (HTTP, SMTP, FTP, and others) are temporarily disabled.
 - Sessions are not dropped when the application parsers are disabled, just the fidelity of the parsing performed on those sessions.
 - Sessions parsed when the application parsers are disabled still have associated network meta (from the network parser).
 - The statistic `/decoder/parsers/stats/blowoff.count` displays the count of all sessions that bypassed application parsers (network parsing is still performed).
 - When session parsing is no longer a potential bottleneck, the application parsers are automatically re-enabled.
 - The assembler session pool should be large enough that it is not forcing sessions.
 - You can determine if sessions are being forced by the statistic `/decoder/stats/asmblr.sessions.forced` (it will be increasing). Also `/decoder/stats/asmblr.sessions` will be within several hundred of `/decoder/config/asmblr.session.pool`.
7. (Optional) If you need to adjust the MTU for capture, add the `snaplen` parameter to `capture.device.params`. Unlike previous releases, the `snaplen` does not need to be rounded up to any specific boundary. The Decoder automatically adjusts the MTU set on the capture interfaces.
8. The following configuration parameters are deprecated and no longer necessary.
 - The `core=` parameter in `capture.device.params`
 - Any configuration files under `/etc/pf_ring` directory
 - Separate `device=` parameters in `capture.device.params`. All multi-interface selection is performed with the **select** command described in Step 5 (b).

Note: An Ethernet device installed post imaging must be added to DPDK if you want to use it as a capture interface. Similarly, it also require configuration if used as a network interface, or for system tools to access it without manual configuration.

Performance Tuning Parameters

By default, the following tunable parameters are disabled. It is recommended that you enable these to achieve high capture rates and consistent performance.

1. You can use the BPF filter to perform fast filtering of the packets. The BPF filter is the fastest way to remove packets from ingest. By dropping the unwanted traffic as early as possible, you can reduce the workload on the Decoder and ensure that essential packets are not dropped. For more information on BPF/PCAP filter, see [Configure Capture Settings](#).
2. You can turn on the Receive Side Scaling (RSS) feature for your network capture interface. The RSS feature splits the traffic coming into the interface in separate queues. It allows each queue to be

handled by a different thread, and therefore run on a different CPU core. This provides more CPU time to execute per-packet operations like evaluating BPF and Network rules. RSS makes handling of higher packet rates easier.

The correct value for RSS will depend on the number of CPUs available on the host. A good starting point for the number of RSS threads depends on how many Cores per CPU socket are present in your Decoder host. Take the number of Cores present on a single CPU, and divide it by 2. For example, on a 12 core processor, you might use up to 6 RSS queues. You can distribute the RSS unequally between interfaces. For example you can assign 4 RSS queues to a busy interface and 2 to a less busy interface.

IMPORTANT: Ensure that a Network Interface on a Decoder host is typically attached to only one CPU socket. Therefore, you must count only the CPU cores of the CPU socket that is attached to the network interfaces.

Setting higher numbers of RSS queues is possible, but there are diminishing returns if the total number of RSS queues spawns more Capture and Assembly threads than there are physical cores on the Decoder host.

For more details on how to configure RSS, see "[Utilizing Receive Side Scaling with DPDK](#)" in [\(Optional\) Data Plane Development Kit Packet Capture](#).

Storage Considerations

Packet Retention Requires Extremely High Sustained Throughput

When capturing at 10G line rates, the storage system holding the packet and meta databases must be capable of sustained write throughput of 1400 MBytes/s. Make a note that many SANs are not capable of achieving these speeds.

Using SAN and Other Storage Configurations

The Decoder allows any storage configuration that can meet the sustained throughput requirement. The standard 8 Gbit FC link to a SAN is not sufficient to store packet data at 10G; in order to use a SAN it may be required to perform aggregation across multiple targets using a software-RAID Scheme. Thus environments using SAN are required to configure connectivity to the SAN using multiple FCs.

Optimize Read/Write Operations When Adding New Storage

A 10G Decoder is optimized to stagger read and write operations across multiple volumes so that the current file being written is on a different volume from the next file that will be written. This allows maximum throughput on the raid volume when reading data from the last file being written while writing the current file on a different volume. However, if volumes are added after a Decoder has been in use, the ability to stagger is limited because one or more volumes are already full so the new volume is the only place new files can be written.

To remedy this situation, an administrator can run a `stagger` command on an existing NetWitness Platform database (packet, log, meta, or session), that has at least two volumes, to stagger the files across all volumes in the most optimal read/write pattern. The major use case is when new storage is added to an existing Decoder and you want to stagger the volumes before restarting capture.

The configuration nodes for this command are the session, meta, and packet databases. Each of these lives under `/database/config`, which is usually a root node. The config nodes for a Decoder are:

- /database/config/packet.dir
- /database/config/meta.dir
- /database/config/session.dir

The *NetWitness Platform Core Database Tuning Guide* has information on how those configurations are formatted.

The `stagger` command is typically only useful for a 10G Decoder and usually just for the packet database. Maximum performance is achieved for storing and retrieving packets when multiple volumes are present. In this scenario, the Decoder always fills the volume with the most free space. When the volumes are roughly the same size, this results in a staggered write pattern, which allows maximum throughput for reading and writing across all volumes. However, this only naturally occurs when multiple packet storage volumes are present at the time the Decoder is first deployed.

A typical use case is adding more storage to an existing Decoder to increase retention. However, when adding storage to an deployment that has already filled the existing volumes with stored packets, the Decoder will naturally fill the new storage with packets before rolling out any packets on the existing storage. This results in a suboptimal read/write pattern because most reads will occur on the same volume that is currently being written to. In a 10G deployment, reads are blocked from the volume when writes are occurring. This does not stop ALL reads on that volume, because the file is buffered in memory before being written, but it does result in suboptimal read performance.

With the `stagger` command, you can add more storage and then have the service naturally stagger the files across ALL volumes (existing and new) so that read performance is optimized.

Caution: This command should only be performed after the storage is mounted and the Decoder configured to use it (for example, after adding the mount point(s) to `packet.dir`).

The downside to this command is it can take some time to stagger and the Decoder should not be capturing during the stagger operation.

Recommended workflow:

1. Add all storage and configure mount points.
2. Add new storage mount points to `packet.dir` (or `session.dir/meta.dir`) and restart service (very important!).
3. Ensure capture is stopped.
4. Run the stagger operation. You must not terminate the connection that initiated the stagger operation until the operation is complete. If you run stagger from NwConsole, run the `timeout 0` command before sending the `stagger` command. This will prevent the normal 30 second command timeout.
5. Start capture after the `stagger` command finishes.

The following are the parameters for the stagger command:

- `type` - The database that will be staggered (session, meta, or packet). Typically only the packet database is useful for staggering, but it is possible to do the session or meta database when multiple volumes are present for those databases. Since the session and meta databases write far less data than the packet database, typically staggering those databases results in less noticeable performance gains.

- `dryRun` - If `true` (the default), will only return a description of the operations that would be performed. If `false`, then the files will actually be moved to an optimal read/write pattern. You MUST pass `false` to actually stagger the files.

Example usage from NwConsole:

```
login <decoder>:50004 <username> <password>
timeout 0
send /database stagger type=packet dryRun=false
```

Note: If you run this command using the RESTful API, pass the additional parameter `expiry=0` to prevent a timeout from the service. You will also need to ensure the HTTP client does not disconnect before the operation completes.

Parsing and Content Considerations

Parsing raw packets at high speeds presents unique challenges. Given the high session and packet rates, parsing efficiency is paramount. A single parser that is inefficient (spends too long examining packets) can slow the whole system down to the point where packets are dropped at the card.

For initial 10G testing, start with only native parsers (except SMB/WebMail). Use the native parsers to establish baseline performance and with little to no packet drops. Do not download any Live content until this has been done and the system is proven to capture without issue at high speeds.

After the system has been operational and running smoothly, Live content should be added very slowly, especially the parsers.

Best Practices

Whether you are updating a currently deployed system or deploying a new system, it is recommended you use the following best practices to minimize risk for packet loss. One caveat is that if you are updating a current 10G deployment but not adding any additional traffic. For example, a current Decoder capturing off a 10G card at 2G sustained should see no difference in performance, unless part of the update also entails adding additional traffic for capture.

- Incorporate baseline parsers (except SMB/Webmail, both of which generally have high CPU utilization) and monitor to ensure little to no packet loss.
- When adding additional parsers, add only one or two parsers at a time.
- Measure performance impact of newly added content, especially during peak traffic periods.
- If drops start occurring when they did not happen before, disable all newly-added parsers and enable just one at a time and measure the impact. This helps pinpoint individual parsers causing detrimental effects on performance. It may be possible to re-factor it to perform better or reduce its feature set to just what is necessary for the customer use case.
- Although lesser performance impacts, feeds should also be reviewed and added in a phased approach to help measure performance impacts.

- Application Rules also tend to have little observable impact, though again, it is best not to add a large number of rules at once without measuring the performance impact.
- If you regularly get a timeout message in the Investigate > Events view, such as The query on channel 192577 was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS), first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or other warnings that may need to be addressed to increase query response time. If there are no messages indicating any specific warnings then try increasing the Core Query Timeout from the default 5 minutes to 10 minutes as described in "View Query and Session Attributes per Role" section of the *System Security and User Management Guide*.

Also, making the recommended configuration changes outlined in the Configuration section will help minimize potential issues.

Use Case 1: 10G Mode - Egress, Deep Packet Inspection

In this setup the goal is to ingest at 10G sustained line rates, perform Deep Packet Inspection (DPI), store metadata, and store raw packets for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- MA content (7 Lua parsers, 1 feed, 1 application rule)
- 4 feeds (alert ids info, nrmalwaredomains, warning, and suspicious)
- 41 application rules
- DNS_verbose_lua (disable DNS)
- fingerprint_javascript_lua
- fingerprint_pdf_lua
- fingerprint_rar_lua
- fingerprint_rtf_lua
- MAIL_lua (disable MAIL)
- SNMP_lua (disable SNMP)
- spectrum_lua
- SSH_lua (disable SSH)
- TLS_lua
- windows_command_shell
- windows_executable

Not Tested

- SMB_lua, native SMB disabled by default
- html_threat

Other

- HTTP_lua reduces the capture rate from >9G to <7G. At just under 5G this parser can be used in place of the native without dropping (in addition to the list above).
- xor_executable pushes parse CPU to 100% and the system can drop significantly due to parse backup.

Aggregation Adjustments Based on Tested Live Content

A 10G Decoder can serve aggregation to a single Concentrator while running at 10G speeds. Deployments using Malware Analysis, Event Stream Analysis, Warehouse Connector, and Reporting Engine are expected to impact performance and can lead to packet loss.

For the tested scenario, the Concentrator aggregates between 45 and 70k sessions per second. The 10G Decoder captures between 40 and 50k sessions per second. With the content identified above, this is about 1.5 to 2 million meta per second. Due to the high volume of session rates, the following configuration changes are recommended:

- Nice aggregation on the Concentrator limits the performance impact on the 10G Decoder. The following command turns on nice aggregation.
`/concentrator/config/aggregate.nice = true`
- Due to the high volume of sessions on the Concentrator, you may consider activating parallel values mode on the Concentrator by setting `/sdk/config/parallel.values` to 16. This improves Investigation performance when the number of sessions per second is greater than 30,000.
- If multiple aggregation streams are necessary, aggregating from the Concentrator instead has less impact on the Decoder.
- Further review for content and parsing is required for deployments where you want to use other NetWitness Platform components (Warehouse Connector, Malware Analysis, ESA, and Reporting Engine).

Configure the Decoder for NDR Mode

The instructions in this section are to configure the Decoder for capture speeds more than 10Gb/sec but less than 40Gb/sec with small amounts of DPI while storing only metadata.

To configure the Decoder:

1. From the **Decoder Explore** view, right-click **Decoder** and select **Properties**. In the properties drop-down menu, select **reconfig** and enter the following parameters:
`op=ndr`
Ensure that the correct parameter is displayed in the output.

2. Add the `update=true` parameter and run **reconfig** again to save the configuration changes. For example, your final parameters will be `op=ndrupdate=true`.
3. From the **Decoder Explore** view, right-click **database** and select **Properties**.
4. In the **Properties** drop-down menu, select **reconfig** and enter the following parameters:
update=1 op=ndr
These parameters adjust the packet database to use very large file sizes and Direct I/O.
5. Perform the following steps to select capture interfaces:
 - a. In the **Decoder Explore** view, right-click on **Decoder** and select **Properties**.
 - b. In the **Properties** drop-down menu, click **select** and run the command to view the available adapters. For example, this might produce output like the list below. The actual list of adapters will vary depending on the hardware available on the Decoder host.

```
1: packet_mmap_,eth0
2: DPDK,0000:41:00.0
3: DPDK,0000:41:00.1
```
 - c. To select a single interface for capture, enter the parameter `adapter=N` where `N` is the capture interface. For example, using the list above you might choose to capture on the first DPDK interface using the `adapter=2` parameter.
 - d. To select more than one interface for capture, enter the parameter `adapter=N,M,...` where `N,M` and so on are the capture interfaces. For example, using the list above you might choose to capture on both DPDK interfaces using the `adapter=2,3` parameter.
6. (Optional) Application parsing is extremely CPU intensive and can cause the Decoder to drop packets. To mitigate application parsing-induced drops, you can set `/decoder/config/assembler.parse.valve` to `true`. These are the results:
 - When session parsing becomes a bottleneck, application parsers (HTTP, SMTP, FTP, and others) are temporarily disabled.
 - Sessions are not dropped when the application parsers are disabled, just the fidelity of the parsing performed on those sessions.
 - Sessions parsed when the application parsers are disabled still have associated network meta (from the network parser).
 - The statistic `/decoder/parsers/stats/blowoff.count` displays the count of all sessions that bypassed application parsers (network parsing is still performed).
 - When session parsing is no longer a potential bottleneck, the application parsers are automatically re-enabled.
 - The assembler session pool should be large enough that it is not forcing sessions.
 - You can determine if sessions are being forced by the statistic `/decoder/stats/assembler.sessions.forced` (it will be increasing). Also `/decoder/stats/assembler.sessions` will be within several hundred of `/decoder/config/assembler.session.pool`.

7. (Optional) If you need to adjust the MTU for capture, add the `snaplen` parameter to `capture.device.params`. Unlike previous releases, the `snaplen` does not need to be rounded up to any specific boundary. The Decoder automatically adjusts the MTU set on the capture interfaces.
8. The following configuration parameters are deprecated and no longer necessary.
 - The `core=` parameter in `capture.device.params`
 - Any configuration files under `/etc/pf_ring` directory
 - Separate `device=` parameters in `capture.device.params`. All multi-interface selection is performed with the **select** command described in Step 5 (b).

Note: An Ethernet device installed post imaging must be added to DPDK if you want to use it as a capture interface. Similarly, it also require configuration if used as a network interface, or for system tools to access it without manual configuration.

Performance Tuning Parameters

By default, the following tunable parameters are disabled. It is recommended that you enable these to achieve high capture rates and consistent performance.

1. You can use the BPF filter to perform fast filtering of the packets. The BPF filter is the fastest way to remove packets from ingest. By dropping traffic that you don't need to retain as early as possible, you can reduce the workload on the Decoder and ensure that essential packets are not dropped. For more information on BPF/PCAP filter, see [Configure Capture Settings](#).
2. You can turn on the Receive Side Scaling (RSS) feature for your network capture interface. The RSS feature splits the traffic coming into the interface in separate queues. It allows each queue to be handled by a different thread, and therefore run on a different CPU core. This provides more CPU time to execute per-packet operations like evaluating BPF and Network rules. RSS makes handling of higher packet rates easier.

The correct value for RSS will depends on the number of CPUs available on the host. A good starting point for the number of RSS threads depends on how many Cores per CPU socket are present in your Decoder host. Take the number of Cores present on a single CPU, and divide it by 2. For example, on a 12 core processor, you might use up to 6 RSS queues. You can distribute the RSS unequally between interfaces. For example you can assign 4 RSS queues to a busy interface and 2 to a less busy interface.

IMPORTANT: Ensure that a Network Interface on a Decoder host is typically attached to only one CPU socket. Therefore, you must count only the CPU cores of the CPU socket that is attached to the network interfaces.

Setting higher numbers of RSS queues is possible, but there are diminishing returns if the total number of RSS queues spawns more Capture and Assembly threads than there are physical cores on the Decoder host.

For more details on how to configure RSS, see "[Utilizing Receive Side Scaling with DPDK](#)" in [\(Optional\) Data Plane Development Kit Packet Capture](#).

Storage Considerations

NDR Mode Assumes No Packet Retention

The base or default configuration for NDR mode starts with all packet writes disabled. So, in this scenario the storage throughput requirements are relatively less. If you choose to turn on packet retention, be aware that it is relatively easy to overwhelm the I/O throughput of most storage solutions with a 40G network feed.

Parsing and Content Considerations

Parsing at Speeds Greater than 10G

As network ingest speeds increase, less CPU time is available to examine each packet or each session. It means that only limited amount of parsing can be performed. Fortunately, you only need to account for the amount of traffic that enters Decoder's parsing sub-system. For example, if you are ingesting 40 Gbit/s but filter out 30 Gbit/s of traffic, you only need to allow 10 Gbit/s of traffic to move through the parsers. At raw 40 Gbit/s speeds, only the "well-behaved" content works. The "well-behaved" content can be defined as:

- Content that only activate when a large, uncommon search token is registered. It includes the Snort rules with a large, uncommon "fast-pattern" content field. For more information on fast-pattern and Snort rules, see [Decoder Snort Detection](#).
- Content that generates a finite number of metas per session.
- Parsers that do not activate on session begin/end events.

The overall maximum throughput of the reassembly and parsing systems on Decoder is about 100,000 sessions (or streams) per second. However, in practical terms, very few downstream analytical or database services can handle such activity. If you find that your session rate is too high, or that your down-stream services cannot handle how many sessions the Decoder is generating, consider filtering out lower value traffic as it is ingested.

Best Practices

The NDR mode has a network rule configured to drop all incoming traffic by default. In general, when capturing above 10Gb/sec, you must limit the amount of traffic analyzed using DPI. These use cases in this document are example guidelines that do not guarantee the actual throughput of the Decoder. As mentioned previously in this document, the actual throughput of the Decoder depends on the following factors:

- **Amount of packets filtered vs. the amount of packets retained**

It is important to filter the traffic that enters the Decoder to minimize packet drops. You can use Network Rules or a BPF Rule to filter the traffic before it enters advanced features such as stream reassembly and DPI. For more information, see [Configure Network Rules](#) and [\(Optional\) Configure System-Level \(BPF\) Packet Filtering](#).

- **Average packet size**

Decoder capturing small packets, for example 64 byte packets, can decrease the overall throughput of

the Decoder. If average packet sizes are small, the line rate of Decoder ingest will fall significantly below the network link speed. It occurs due to the amount of per-packet overhead between each Ethernet frame.

- **Shared system resources**

If the Decoder host is also hosting other applications or add-on features that manipulate data and perform other analysis, it can lower the overall throughput of the Decoder. Even if these add-on applications are not utilizing the CPU, these will utilize the shared resources like memory bandwidth and CPU interconnect bandwidth.

Whether you are updating a currently deployed system or deploying a new system, it is recommended you use the following best practices to minimize risk for packet loss.

- Incorporate baseline parsers (except SMB/Webmail, both of which generally have high CPU utilization) and monitor to ensure little to no packet loss.
- When adding additional parsers, add only one or two parsers at a time.
- Measure performance impact of newly added content, especially during peak traffic periods.
- If drops start occurring when they did not happen before, disable all newly-added parsers and enable just one at a time and measure the impact. This helps pinpoint individual parsers causing detrimental effects on performance. It may be possible to re-factor it to perform better or reduce its feature set to just what is necessary for the customer use case.
- If you regularly get a timeout message in the Investigate > Events view, such as `The query on channel 192577 was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS)`, first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or other warnings that may need to be addressed to increase query response time. If there are no messages indicating any specific warnings then try increasing the Core Query Timeout from the default 5 minutes to 10 minutes as described in "View Query and Session Attributes per Role" section of the *System Security and User Management Guide*.

Use Case 1: NDR Mode - Egress, General Purpose

Generate NetFlow Style Meta Only + Small Subset of Snort Rules + All Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- The **NETWORK** native parsers
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.
- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see [Configure Capture Settings](#).

Use Case 2: NDR Mode - Egress, Data Exfiltration

Generate NetFlow Style Meta Only + Exfiltration Focused Specific Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- Native parsers - HTTP, HTTPS(SSL), SMTP, DNS, FTP, SFTP/SSH, and VNC
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.
- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see [Configure Capture Settings](#).

Use Case 3: NDR Mode - Lateral Movement

Generate NetFlow Style Meta Only + Small Subset of Snort Rules + Specific Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- Native parsers - Kerberos, SMB, VNC, and SFTP/SSH
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

- Disable the Snort rule related to SMB traffic (for example, `M.HackTool.SMB.Impacket-Obfuscation.[Service Names]`). It is recommended because large regex parameters can cause Snort parser to utilize more than 50% CPU.
- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.
- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see [Configure Capture Settings](#).


Configure a Log Decoder to Accept Protobuf

There are occasions when you want to analyze log files that are in protobuf (Protocol Buffer) format. You can configure a Log Decoder with a Log Collector service to accept logs in protobuf (Protocol Buffer) format.

To import a log file to a Log Decoder:


1. Go to  (Admin) > Services.
2. Select a Log Decoder with a Log Collector service in the **Service** list, and select  > **View** > **Explore**.

The Explore view for the Log Decoder -Log Collector is displayed.



3. Navigate to `event-processors/logdecoder/destinations/logdecoder/consumer/processors/tcpconnector/config`

Your screen should look similar to the following.



4. For the **send-protobuf** field, select **false**, and change the value to **true**.
5. Navigate to event-processors/logdecoder/destinations/logdecoder/consumer/processors/tcpconnector/config/connector/channel/tcp and change the **port** value to **50202**.
6. Navigate to event-processors/logdecoder/destinations/logdecoder/consumer/processors/tcpconnector/config/connector/event and change the following parameters:
 - Clear the **delimiter** field
 - Change **format** to **%text%**

Configure Session Split Timeouts

The default behavior of the Decoder is to automatically end sessions that exceed a configured size or have been inactive for a period of time. When the session is ended due to timeout, any subsequent packets received in that session appear to be stored in a new session. You can mitigate the effect of session splitting due to long periods of inactivity between packets using this procedure.

When a Decoder session exceeds a configured size (32MB by default, the `/decoder/config/assembler.max.size`) or has been inactive for a period of time, the session is split. NetWitness Platform has the previous packet and the next packet and can propagate session state from the initial session fragment to the subsequent session fragment.

Each session fragment is annotated (`session.split` meta) such that it can be identified and associated with other fragments from the actual network session. Directionality as determined by the initial session reduces the occurrence of fragments having reversed directionality.

If there is a gap in time between packets large enough that there are no longer any packets for the session in memory, the session is removed from the Decoder. If a subsequent packet shows up after this occurs, a new session is created with no context to the preceding session. The issue is the inability to continue a session when we encounter a gap between packets of a session that is larger than the packets we are buffering (based upon available memory and timeout configurations). Once the last packet of a session is removed from memory, the session is also removed, and with it the necessary context for ensuring consistent directionality.

There are two timeout settings in a Network Decoder, `/decoder/config/assembler.timeout.session` and `assembler.timeout.packet`. Both default to 60 seconds. The setting `assembler.timeout.session` controls how long a session lives in Assembler without receiving another packet. The setting `assembler.timeout.packet` controls how long a session waits before getting parsed. If the session is kicked out of Assembler before this timeout, then it automatically goes to parsing.

The session timeout is the number of seconds since the last packet was added to that session. Therefore, this timeout resets on every packet added to that session. The packet timeout is the number of seconds since the very first packet for that session was added (in other words, the packet that created the session). This is never reset and once the timeout expires, the session is parsed.

The important point is a session can be parsed but still remain in Assembler. A session in Assembler can still have packets added to it, even if it has already been parsed. Packets added after the session is parsed will never be seen by parsers, but they will be attached to the session and can be viewed by a subsequent `/sdk content` or `/sdk packets` call.

After a session is parsed, the session AND its metadata are written to disk. At this point, they can be aggregated and "seen" by `sdh` commands. Packets are written in order of capture and are not reordered by what session they belong to. Nor are they necessarily written when the session and meta data are written.

You can disable both timeout nodes, `/decoder/config/assembler.timeout.session` and `assembler.timeout.packet`, by setting them to zero in the Services Explore view.


If both timeouts are disabled, the sessions are still split due to time or size expiration. However, the Decoder keeps track of the network stream for as long as it has sufficient memory. Thus, when more packets arrive on the same network stream, the Decoder adds `split` meta items to the subsequent sessions. Using a combination of the `split` metadata and the stream key, it is possible to reconstruct the network stream from the multiple sessions.

The length of time for which sessions are tracked is limited by the number of session pool entries available on the Decoder, and therefore the actual time window varies according to the rate at which new sessions are added. If new sessions are added at a high rate, the size of the time window decreases. The size of the pool is set using the configuration entry `/decoder/config/assembler.session.pool`, which sets the maximum number of sessions that will be tracked at a time.

The `/decoder/stats/assembler.timespan` statistic allows you to see when the Decoder is no longer tracking session splits because the ingest rate is too high and the Decoder does not have enough memory to track. This statistic shows the number of seconds tracked within the session table, which is the effective time window in which the Decoder can link together sessions. Under normal operation this statistic matches the value of `/decoder/config/assembler.timeout.session`, but when running in Time Split mode, the `/decoder/stats/assembler.timespan` statistic grows or shrinks depending on the ingest rate.

To configure Time Split mode, set the following configuration parameters and restart the Decoder:

1. In the  (Admin) > **Services** view, select the Decoder service and  > **View** > **Explore**.
2. In the Services Explore view select **decoder** > **config**.



The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. Below this is a secondary navigation bar with 'HOSTS', 'SERVICES', 'EVENT SOURCES', 'ENDPOINT SOURCES', 'HEALTH & WELLNESS', 'SYSTEM', and 'SECURITY'. The 'SERVICES' section is active, showing a list of services: 'Decoder', 'connections', 'database', 'decoder', 'config', 'devices', 'parsers', 'stats', 'deviceappliance', 'index', 'logs', 'rest', 'sdk', 'services', 'storedproc', 'sys', and 'users'. The 'config' service is selected, and the 'Explore' view is open. The main panel displays a table of configuration parameters for the Decoder service.

Parameter	Value
<code>/decoder/config/adapt.meta.enabled</code>	on
<code>/decoder/config/aggregate.precache</code>	on
<code>/decoder/config/assembler.parse.valve</code>	false
<code>/decoder/config/assembler.pool.ratio</code>	70
<code>/decoder/config/assembler.session.flush</code>	1
<code>/decoder/config/assembler.session.pool</code>	2000
<code>/decoder/config/assembler.size.max</code>	32 MB
<code>/decoder/config/assembler.size.min</code>	0
<code>/decoder/config/assembler.tcp.close</code>	false
<code>/decoder/config/assembler.timeout.packet</code>	60
<code>/decoder/config/assembler.timeout.session</code>	60
<code>/decoder/config/assembler.voting.weights</code>	first=1 size=1 port=1 octet=1 routable=1
<code>/decoder/config/capture.autostart</code>	on
<code>/decoder/config/capture.buffer.size</code>	8 MB
<code>/decoder/config/capture.device.params</code>	
<code>/decoder/config/capture.selected</code>	packet_mmap_eth0 (bpf)
<code>/decoder/config/export.cache.expire</code>	60
<code>/decoder/config/export.packet.enabled</code>	no
<code>/decoder/config/export.packet.local.path</code>	/var/netwitness/decoder/export/packet

3. Click in the **Value** column next to the parameter and set these two parameters :
`/decoder/config/assembler.session.flush = 0`
`/decoder/config/assembler.timeout.session = 0`
4. To see when the Decoder is no longer tracking session splits because the ingest rate is too high and the Decoder does not have enough memory to track, view the `/decoder/stats/assembler.timespan` statistic, in the Services Explore view select **decoder** >

Configure Syslog Forwarding to Destination

In addition to collecting Syslog messages, you can configure the Log Decoder to forward Syslog messages to another Syslog receiver.

Note: No configuration is necessary to collect Syslog messages on the Log Decoder.

NetWitness Platform forwards Syslog messages after it has parsed the messages and before it writes the messages to the Log Decoder.

Note: You must configure Syslog Forwarding using the steps defined in this topic using the **Explore** view.


The Log Decoder must be in the **Started** state before you can configure Syslog Forwarding. To configure Syslog Forwarding:

1. Configure Log Decoder application layer rules (Application rules) to tag Syslog messages with metadata that instructs NetWitness Platform to forward the messages:
 - a. In the **Services** view, select a Log Decoder, and in the Actions column, select   > **View** > **Explore**.
 - b. Go to the `/decoder/config/rules/application` node, right-click **application**, and click **Properties**.
 - c. In the **Properties** view, specify the **add** command with the following parameters:


```
rule=<query> name=<name>
```

Example 1: `rule=* name=receiver1`

Example 2: `rule="device.type='winevent_nic'" name=receiver)`
 - d. Click **Send**.



NetWitness Platform creates the `name=receiver1 rule=* order=<n>` rule. NetWitness Platform inserts the order number (for example, `order=49`) based on when you set up the rule.

0049

`rule=* name=receiver1 order=49`

- e. Go to the `/decoder/config/rules/application` node and click the `name=receiver1 rule=* order=49` rule.

- f. Add **alert forward** parameters to the rule parameters.

```
rule=* name=receiver1 order=49 alert forward
```

or

```
rule=* name=receiver1 forward
```

All other rule parameters have the same meaning as they do in other application rules.

The following Application rule example selects all logs with the `*` rule. It creates a transient alert meta with the value "receiver1" and tags the entire log for forwarding it to the syslog forwarding destination. You can define as many different forwarding rules as you need with the same name or unique names.

2. Define Syslog forwarding destinations and enable forwarding.

- a. In the **Services** view, select a Log Decoder, and   > **View** > **Explore**.

- b. Syslog forwarding destinations are defined in the configuration node.

```
/decoder/config/logs.forwarding.destination
```

This configuration node contains one or more name/value pairs. The name corresponds to the name parameter in the application rule that you used to tag logs with forwarding meta. The value is a colon-separated triple of transport, host, and port followed by an optional formatting parameter.

```
name=(udp|tcp|tls):host:port[: (retainsource|rfc3164)]
```

The first parameter indicates the transport protocol and must be one of `udp`, `tcp`, or `tls`. Specifying `udp` will forward logs via RFC 3164 / RFC 5426 UDP syslog protocol. Specifying `tcp` will forward logs via a TCP connection with RFC 6587 framing. Specifying `tls` will forward logs in accordance with RFC 5425.

The host is an IPv4 address, IPv6 address, or host name.

The port is the port to which the logs are sent. This is typically port 514 for UDP syslog, and 6514 for TLS connections. There is no standard port assignment for syslog over TCP.

Optionally, `retainsource` or `rfc3164` can be specified at the end of the destination string to indicate that additional formatting and information should be included with each log forwarded. Specifying `retainsource` will include z-connector headers at the beginning of the log based and will be populated by the `time`, `device.(ip|ipv6|host)`, and `lc.cid` meta and is best used for forwarding to other log decoders. The `rfc3164` option will prepend a valid RFC3164 header to all events forwarded constructed of the `syslog.pri`, `time`, and `device.(ip|ipv6|host)` meta. In both cases, the original log text is unmodified.

Example forwarding destination:

```
gears=tls:gears.netwitness.local:6514
```

Example forwarding over tcp to blackout on port 514 with z-connector headers:

```
fwdrule=tcp:blackout.netwitness.local:514:retainsource
```

In the `/decoder/config/logs.forwarding.destination` parameter, specify the destination. For example:

TLS Connections: `receiver1=tls:receiver1.netwitness.local:6514`

UDP Connections: `receiver1=udp:receiver1.netwitness.local:514`

TCP Connections: `receiver1=tcp:receiver1.netwitness.local:514`

<code>logs.forwarding.destination</code>	<code>receiver1=tcp:10.31.244.44:514 receiver2=tcp:10.31.244.46:514 receiver3=tcp:10.31.244.48:514</code>

Note:

You can configure:

- Multiple rules to forward logs to the same destination.
- Multiple rules to forward logs to multiple destination.

For TLS connections, the certificate of the forwarding destination must be validated. The certificate authority that signed the destination's certificate must be present in the Log Decoder's CA trust store and the certificate must reside on the destination or Syslog receiver. Refer to "Configure Certificates" in the *Log Collection Configuration Guide* for information about manipulating the Log Decoder's CA trust store. (Go to the [Master Table of Contents](#) to find all RSA NetWitness Platform 11.x documents.)

- c. In the `/decoder/config/logs.forwarding.enabled` parameter, specify **true**.

<code>logs.forwarding.enabled</code>	<code>true</code>

Configure Transaction Handling on a Decoder

Administrators can configure a Decoder to subdivide incoming sessions into smaller transaction sessions when using Lua parsers designed to create transactions. The feature allows analysts to perform analytics on the split sessions in downstream services such as Investigate.

Caution: Use caution when enabling this feature when you are above the standard ingest rate supported by the Network Decoder as it may start to cause backups with aggregation and dropped packets.

Transaction Handling

The Decoder service configuration node has a new parameter for configuration of transaction handling: `/decoder/parsers/config/parser.transaction.mode`. This node controls the behavior of the Decoder when a parser defines a transaction within a network session.

The values for `parser.transaction.mode` correspond to the operating modes:

- `off` (transactions off)
- `meta` (transactions represented as meta items)
- `split` (transactions split sessions)

Transactions Off

When transactions mode is off, any application-level transactions created by parsers are ignored, and nothing is stored in the collection to represent the transaction.

Transactions Represented as Meta Items

In this mode of operation, when a parser generates an application-level transaction, a new meta item of type `{{trans}}` is added to the session in which the transaction occurred. The `{{trans}}` meta item contains a list of other meta items that constitute the transaction.

Transactions Split Sessions

In this mode of operation, when a parser generates an application-level transaction, the session is split. The session splitting is accomplished by:

1. A new session item is created.
2. Network meta items are copied from the parsed session into the new session.
3. Meta items marked in the transaction are moved from the original session to the new session.

The following meta items are duplicated into the split session from the session that was parsed:

- `time`
- `medium`
- `eth.src`
- `eth.dst`
- `eth.type`
- `ip.proto`
- `ip.src`
- `ip.dst`
- `ipv6.src`
- `ipv6.dst`
- `ip.proto`
- `port.src`
- `port.dst`
- `tcp.flags`
- `udp.srcport`
- `udp.dstport`
- `service`
- `udp.srcport`
- `udp.srcport`
- `tls.premaster`

Configure Data Export

This exports meta data and raw logs from the Decoder or Log Decoder and converts it to Logstash JSON object. To enable data export you must configure the export connector through Logstash.

To configure the data export:


1. Go to  (**Admin**) > **Services** and select a Decoder or Log Decoder.
2. Select  > **View** > **Config**.
3. In the Configuration view, select the **Data Export** tab.
4. In the Data Export list, click the name of the host.

The Event Source tab is displayed with export connector typespec selected.


For more information on configuring export connector, see [Configure Logstash Event Sources in NetWitness Platform](#).

Decrypt Incoming Packets

Administrators can configure a Network Decoder to decrypt incoming network packets. This capability is currently supported with encrypted communications to managed servers using TLS 1.2 or earlier and certain cipher suites, which are described in [Supported Cipher Suites](#). The following figure illustrates the supported configuration of a Network Decoder deployed as a passive network device where it captures network traffic for managed servers. The Network Decoder is able to decrypt the traffic by using private keys uploaded to it from the managed servers.



As the network traffic is collected and passed through the system, the encrypted packets are stored as captured even if the system has been able to decrypt them. As the Decoder parses the encrypted traffic, it uses any private keys that have been uploaded, and determines if the network traffic can be decrypted. If the Decoder is able to decrypt the traffic, it generates an additional meta key, defined as `tls.premaster`, which contains the decryption key for that particular session. Whether the traffic is decrypted or not, the metadata service is tagged as 443 or SSL/TLS traffic, as the Decoder is port agnostic, even though it is marked with the default port number. The following figure shows the basic progression through the NetWitness platform as the packets are being decrypted and the session metadata is generated.



Follow the steps below to configure the Network Decoder to natively decrypt this network traffic.

- [Step 1: Validate That The Network Decoder Captures Encrypted Traffic](#)
- [Step 2: Obtain Private Keys from Managed Servers](#)
- [Step 3: Validate That The Private Key Cipher Suite is Supported](#)
- [Step 4: Confirm HTTPS Parser is Enabled on Decoders](#)
- [Step 5: Upload the Supported Private Keys to Decoders](#)
- [Step 6: Reload HTTPS Parser on Decoders](#)
- [Step 7: View Decrypted Sessions](#)

Note: Currently, you cannot review the decrypted payloads directly in the in the user interface. However, you can review them using the RESTful API interface and the NwConsole. However, the Decoders being able to decrypt traffic enables them to detect and generate potentially malicious communications that would otherwise not be detected.

See [Troubleshooting](#) for resolving any issues you might encounter while performing these steps.

Step 1: Validate That The Network Decoder Captures Encrypted Traffic

First, make sure the Decoder is actually capturing the SSL/TLS traffic from the managed servers that you want to use to gain visibility into the encrypted sessions. You can do this by searching in the Investigate views (Navigate or Events) using the IP addresses of the servers as the destination IP addresses, and the service as SSL/TLS.

For example, if you are searching for HTTPS traffic to 192.168.1.2 and 192.168.1.3, you can run the following query:

```
service = 443 && ip.dst = 192.168.1.2, 192.168.1.3
```

Conduct this search on the Concentrator that aggregates metadata from the Decoder that is expected to capture this network segment, or on a Broker that may have access to all collected metadata.

(Version 11.6 or later) In case of simultaneous ingestion of the encrypted and decrypted traffic streams in the Decoder, ensure that multi-thread assembly is enabled. This is to prevent mixing of packets from different sessions by single thread assembly. For more information, see the following topics:

- [Perform Simultaneous Ingestion of the Encrypted and Decrypted Traffic Streams to Decoder](#)
- [\(Optional\) Multiple Adapter Packet Capture](#)

Step 2: Obtain Private Keys from Managed Servers

Private keys are the asymmetric keys generated by services that accept TLS traffic. They are normally stored in PEM files. These keys are used during the TLS handshake to encrypt the premaster symmetric key that is used for the rest of the payload encryption.

Note: The method to export private keys from web servers (from which to decrypt traffic) varies depending on the product. Refer to the product documentation for instructions for exporting private keys.

Step 3: Validate That The Private Key Cipher Suite is Supported

As a passive collection device, the Network Decoder can only decrypt ciphers that use the RSA key exchange. The list of supported keys, and whether the keys are FIPS compliant, are listed in [Supported Cipher Suites](#) and [Unsupported Cipher Suites](#). If you are expecting to decrypt less-secure channels of communication, you might need to disable FIPS. However, this is a system change to the entire Decoder, not just the parsing of captured traffic, so take caution before disabling FIPS. If you require more-secure cipher suites for the managed server communications, search for "SSL decrypt" in our third party partner integrations section at RSA Ready Community (<https://community.rsa.com/community/products/rsa-ready>).

Note: If FIPS is enabled, the list of ciphers for decryption is restricted to only those that are FIPS-compliant.

Supported Cipher Suites

The following table lists which cipher-suites are supported using private keys.

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_RSA_WITH_AES_256_GCM_SHA384	AES256-GCM-SHA384	[0x9d]	TLSv1.2	Kx=RSA	Compliant
TLS_RSA_WITH_AES_256_CBC_SHA256	AES256-SHA256	[0x3d]	TLSv1.2	Kx=RSA	Compliant

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_RSA_WITH_AES_256_CBC_SHA	AES256-SHA	[0x35]	SSLv3	Kx=RSA	Compliant
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	CAMELLIA256-SHA	[0x84]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_3DES_EDE_CBC_SHA	DES-CBC3-SHA	[0x0a]	SSLv3	Kx=RSA	Compliant
TLS_RSA_WITH_AES_128_GCM_SHA256	AES128-GCM-SHA256	[0x9c]	TLSv1.2	Kx=RSA	Compliant
TLS_RSA_WITH_AES_128_CBC_SHA256	AES128-SHA256	[0x3c]	TLSv1.2	Kx=RSA	Compliant
TLS_RSA_WITH_AES_128_CBC_SHA	AES128-SHA	[0x2f]	SSLv3	Kx=RSA	Compliant
TLS_RSA_WITH_SEED_CBC_SHA	SEED-SHA	[0x96]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	CAMELLIA128-SHA	[0x41]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_IDEA_CBC_SHA	IDEA-CBC-SHA	[0x07]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_RC4_128_SHA	RC4-SHA	[0x05]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_DES_CBC_SHA	DES-CBC-SHA	[0x09]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_EXPORT_WITH_DES40_CBC_SHA	EXP-DES-CBC-SHA	[0x08]	SSLv3	Kx=DES	Non-Compliant

Unsupported Cipher Suites

The following table lists which cipher-suites are not supported using private keys.

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	ECDHE-RSA-AES256-GCM-SHA384	[0xc030]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	ECDHE-ECDSA-AES256-GCM-SHA384	[0xc02c]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	ECDHE-RSA-AES256-SHA384	[0xc028]	TLSv1.2	Kx=ECDH	Compliant

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	ECDHE-ECDSA-AES256-SHA384	[0xc024]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	ECDHE-RSA-AES256-SHA	[0xc014]	SSLv3	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	ECDHE-ECDSA-AES256-SHA	[0xc00a]	SSLv3	Kx=ECDH	Compliant
TLS_DHE_DSS_WITH_AES_256_GCM_SHA384	DHE-DSS-AES256-GCM-SHA384	[0xa3]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	DHE-RSA-AES256-GCM-SHA384	[0x9f]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	DHE-RSA-AES256-SHA256	[0x6b]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256	DHE-DSS-AES256-SHA256	[0x6a]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_256_CBC_SHA	DHE-RSA-AES256-SHA	[0x39]	SSLv3	Kx=DH	Compliant
TLS_DHE_DSS_WITH_AES_256_CBC_SHA	DHE-DSS-AES256-SHA	[0x38]	SSLv3	Kx=DH	Compliant
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	DHE-RSA-CAMELLIA256-SHA	[0x88]	SSLv3	Kx=DH	Non-Compliant
TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA	DHE-DSS-CAMELLIA256-SHA	[0x87]	SSLv3	Kx=DH	Non-Compliant
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384	ECDH-RSA-AES256-GCM-SHA384	[0xc032]	SSLv3	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384	ECDH-ECDSA-AES256-GCM-SHA384	[0xc02e]	TLSv1.2	Kx=ECDH/ECDSA	Compliant
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384	ECDH-RSA-AES256-SHA384	[0xc02a]	TLSv1.2	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384	ECDH-ECDSA-AES256-SHA384	[0xc026]	TLSv1.2	Kx=ECDH/ECDSA	Compliant
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	ECDH-RSA-AES256-SHA	[0xc00f]	SSLv3	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	ECDH-ECDSA-AES256-SHA	[0xc005]	SSLv3	Kx=ECDH/ECDSA	Compliant

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	ECDHE-RSA-DES-CBC3-SHA	[0xc012]	SSLv3	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	ECDHE-ECDSA-DES-CBC3-SHA	[0xc008]	SSLv3	Kx=ECDH	Compliant
TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	EDH-RSA-DES-CBC3-SHA	[0x16]	SSLv3	Kx=DH	Compliant
TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	EDH-DSS-DES-CBC3-SHA	[0x13]	SSLv3	Kx=DH	Compliant
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	ECDH-RSA-DES-CBC3-SHA	[0xc00d]	SSLv3	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	ECDH-ECDSA-DES-CBC3-SHA	[0xc003]	SSLv3	Kx=ECDH/ECDSA	Compliant
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	ECDHE-RSA-AES128-GCM-SHA256	[0xc02f]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	ECDHE-ECDSA-AES128-GCM-SHA256	[0xc02b]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	ECDHE-RSA-AES128-SHA256	[0xc027]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	ECDHE-ECDSA-AES128-SHA256	[0xc023]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	ECDHE-RSA-AES128-SHA	[0xc013]	SSLv3	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	ECDHE-ECDSA-AES128-SHA	[0xc009]	SSLv3	Kx=ECDH	Compliant
TLS_DHE_DSS_WITH_AES_128_GCM_SHA256	DHE-DSS-AES128-GCM-SHA256	[0xa2]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	DHE-RSA-AES128-GCM-SHA256	[0x9e]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	DHE-RSA-AES128-SHA256	[0x67]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_DSS_WITH_AES_128_CBC_SHA256	DHE-DSS-AES128-SHA256	[0x40]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_128_CBC_SHA	DHE-RSA-AES128-SHA	[0x33]	SSLv3	Kx=DH	Compliant


Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_DHE_DSS_WITH_AES_128_CBC_SHA	DHE-DSS-AES128-SHA	[0x32]	SSLv3	Kx=DH	Compliant
TLS_DHE_RSA_WITH_SEED_CBC_SHA	DHE-RSA-SEED-SHA	[0x9a]	SSLv3	Kx=DH	Non-Compliant
TLS_DHE_DSS_WITH_SEED_CBC_SHA	DHE-DSS-SEED-SHA	[0x99]	SSLv3	Kx=DH	Non-Compliant
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA	DHE-RSA-CAMELLIA128-SHA	[0x45]	SSLv3	Kx=DH	Non-Compliant
TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA	DHE-DSS-CAMELLIA128-SHA	[0x44]	SSLv3	Kx=DH	Non-Compliant
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256	ECDH-RSA-AES128-GCM-SHA256	[0xc031]	TLSv1.2	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256	ECDH-ECDSA-AES128-GCM-SHA256	[0xc02d]	TLSv1.2	Kx=ECDH/ECDSA	Compliant
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256	ECDH-RSA-AES128-SHA256	[0xc029]	TLSv1.2	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256	ECDH-ECDSA-AES128-SHA256	[0xc025]	TLSv1.2	Kx=ECDH/ECDSA	Compliant
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	ECDH-RSA-AES128-SHA	[0xc00e]	SSLv3	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	ECDH-ECDSA-AES128-SHA	[0xc004]	SSLv3	Kx=ECDH/ECDSA	Compliant
TLS_ECDHE_RSA_WITH_RC4_128_SHA	ECDHE-RSA-RC4-SHA	[0xc011]	SSLv3	Kx=ECDH	Non-Compliant
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	ECDHE-ECDSA-RC4-SHA	[0xc007]	SSLv3	Kx=ECDH	Non-Compliant
TLS_ECDH_RSA_WITH_RC4_128_SHA	ECDH-RSA-RC4-SHA	[0xc00c]	SSLv3	Kx=ECDH/RSA	Non-Compliant
TLS_ECDH_ECDSA_WITH_RC4_128_SHA	ECDH-ECDSA-RC4-SHA	[0xc002]	SSLv3	Kx=ECDH/ECDSA	Non-Compliant
TLS_DHE_RSA_WITH_DES_CBC_SHA	EDH-RSA-DES-CBC-SHA	[0x15]	SSLv3	Kx=RSA	Non-Compliant
TLS_DHE_DSS_WITH_DES_CBC_SHA	EDH-DSS-DES-CBC-SHA	[0x12]	SSLv3	Kx=DSS	Non-Compliant
TLS_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA	EXP-EDH-RSA-DES-CBC-SHA	[0x14]	SSLv3	Kx=DSS	Non-Compliant

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA	EXP-EDH-DSS-DES-CBC-SHA	[0x11]	SSLv3	Kx=DSS	Non-Compliant

Step 4: Confirm HTTPS Parser is Enabled on Decoders

The decryption process requires the native HTTPS parser to be enabled to examine the TLS handshake and to compare the uploaded encryption keys to determine if they decrypt the traffic. This can be

validated in  (Admin) > **Services** > **Decoder** > **Config** > Parsers Configuration panel as shown in the following figure.



The screenshot shows the 'Decoder Configuration' interface. On the left, there are sections for 'System Configuration' and 'Decoder Configuration'. The 'Parsers Configuration' panel on the right is active, showing a list of parsers and their status. The 'HTTPS' parser is highlighted with a red box, and its sub-parsers are also listed and enabled.

Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☒
SSL Port	56004
Stat Update Interval	1000
Threads	20

Name	Config Value
Adapter	Berkeley Packet Filter
Capture Interface Selected	packet_mmap_eth0 (bpf)
Cache	Cache Directory: /var/netwitness/decoder/cache
Cache Size	4 GB
Capture Settings	Assembler Maximum Size: 32 MB

Name	Config Value
HTTPS	Enabled
cert.sha256	Enabled
cert.thumbprint	Enabled
client	Enabled
crypto	Enabled
ja3	Enabled
ja3.orig	Enabled
ja3s	Enabled
ja3s.orig	Enabled
tls.cipherlen	Enabled
tls.extensionlen	Enabled
tls.premaster	Enabled
IRC	Enabled
JSNormalize	Enabled
MAIL	Enabled
NETBIOS	Enabled

Step 5: Upload the Supported Private Keys to Decoders

The `sslKeys` command provides a way to upload premaster or private keys to the Decoder, so that captured encrypted packets that match the keys can be decrypted before parsing. Administrators configure the Decoder by entering the `sslKeys` command using the NwConsole command line interface or the Decoder RESTful interface. After upload, you must issue a parser reload command so that the newly installed key becomes visible to the HTTPS parser. Now, all TLS handshakes that use that private key will be able to be decrypted by the Decoder.

For information about how to use the NwConsole command line interface, see the *NwConsole User Guide for RSA NetWitness Platform*. For information about how to use the RESTful interface, see the *RESTful API User Guide for RSA NetWitness Platform*. Go to the [Master Table of Contents](#) to find all RSA NetWitness Platform 11.x documents.

The screenshot shows the NwConsole RESTful interface. At the top, there's a navigation bar with 'SERVICES' and a dropdown menu showing 'storedproc (*)', 'sys (*)', and 'users (*)'. Below this, the 'Properties for /decoder' section has a dropdown menu set to 'sslKeys' and a 'Parameters:' input field with a 'Send' button. The 'Message Help' section contains the following text: 'sslKeys: Push SSL crypto information to enable SSL decryption of a session's packets prior to parsing', 'security.roles: decoder.manage', and 'parameters: clear - <bool, optional> Clears all existing keys from storage. Cannot be used with any other parameters.', 'maxKeys - <uint32, optional> Sets the total number of keys that can be held in memory before aging out begins. Cannot be used with any other parameters.', and 'random - <string, optional> Adds the random that identifies the session key exchange.' Below this is an 'Output (or command manual help)' section. It contains a paragraph explaining the 'premaster' key and how to generate it using Chrome or Firefox. It then provides a sample NwConsole script: 'login <decoder>:50004 <username> <password>', 'send /decoder sslKeys --file-data=SSLKeys.txt', and a curl command: 'curl -u "<username>:<password>" -H "Content-Type: application/octet-stream" --data-binary @"/path/SSLKeys.txt" -X POST "http://<host>:50004/decoder/sslKeys"'. It also mentions that symmetric keys are stored in memory and there is a limit to how many can be stored. The 'Private Keys or PEM files' section is partially visible at the bottom.

The RESTful interface form at the path `/decoder/sslkeys` allows uploading a single PEM-encoded private key, a single file containing multiple private keys concatenated together, or a single file of multiple premaster keys.

SSL Decryption Key Upload

Use this form to upload SSL decryption key information.
Types of uploads supported:

- PEM-encoded private keys, optionally with many can be concatenated in the same file.
- Premaster keys in NSS Key Log Format

You can use the `sslKeys` message on the `/decoder` folder to view or modify the current set of decryption keys.

Choose up to 3 files to upload.

Upload file 1:  AES256-GC...HA384.pem

Upload file 2:  SSLKeysTLS11.txt

Upload file 3:  super.pem




[Back to Root Folder](#)

Upload Multiple Premaster and Private Keys

Note: As of 11.5, uploading premaster keys is not officially supported yet for real-time capture, but can be used to validate that decryption is functioning properly on the Decoder.

You can use the RESTful interface form to facilitate uploading of multiple keys, both premaster and private at the same time.

1. Open the RESTful API in your browser and navigate to this path on the Decoder that you want to configure: `/decoder/sslkeys`.



The screenshot shows a web browser window with the address bar displaying `localhost:50104/decoder/sslkeys`. The page title is "SSL Decryption Key Upload". The main content area has a light gray background and contains the following text and form elements:

- SSL Decryption Key Upload**
- Use this form to upload SSL decryption key information.
- Types of uploads supported:
 - PEM-encoded private keys, optionally with many keys concatenated in the same file.
 - Premaster keys in NSS Key Log Format
- You can use the `sslKeys` message on the `/decoder` folder to view or modify the current set of decryption keys.
- Choose up to 3 files to upload.
- Upload file 1: no file selected
- Upload file 2: no file selected
- Upload file 3: no file selected
-
- [Back to Root Folder](#)

2. Next to **Upload File 1**, click **Choose File** and locate the premaster key file or PEM file that you want to upload on your local file system.
3. (Optional) Repeat for **Upload File 2** and **Upload File 3**.

SSL Decryption Key Upload

Use this form to upload SSL decryption key information.
Types of uploads supported:

- PEM-encoded private keys, optionally with many can be concatenated in the same file.
- Premaster keys in NSS Key Log Format

You can use the sslKeys message on the /decoder folder to view or modify the current set of decryption keys.

Choose up to 3 files to upload.

Upload file 1: AES256-GC...HA384.pem


Upload file 2: SSLKeySTLS11.txt

Upload file 3: super.pem

[Back to Root Folder](#)

4. Click **Upload**.

The files are uploaded to the Decoder and results are displayed in the form.



Parameters for Managing Keys

The `sslKeys` command has several parameters for managing premaster and private keys. This is the full list of parameters:

Parameter	Description
<code>clear</code>	Removes all premaster keys from memory. Does not delete any PEM files installed on the system.
<code>maxKeys</code>	Changes the maximum number of premaster keys that are stored in memory.
<code>listPems</code>	Returns a list of all installed private key PEM files.
<code>deletePem</code>	Deletes the named PEM file from the file system. You can pass this parameter more than once to remove multiple files.
<code>random</code>	The random hash used to identify the premaster key.
<code>premaster</code>	The premaster key that will be installed for the previous <code>random</code> parameter. They must show up in pairs and <code>random</code> must be first.

Return Values

Most `sslKeys` commands return name/value pairs of statistics about the premaster keys in memory. The statistics are listed in the following table.

Name	Description
added	The number of premaster keys just added during this command.
total	The total number of premaster keys loaded in memory.
agedOut	The total number of premaster keys that were removed during this command; this is not a lifetime stat.
maxKeys	The current maximum allowed premaster keys

Encryption Keys

The `sslKeys` command accepts two types of encryption keys:

- Premaster key: the symmetric key used in the TLS payload stream for encryption and decryption.
- Private key: the asymmetric private key used during the TLS handshake that encrypts the premaster.

Premaster Key

The premaster key is generated randomly and is ephemeral for the life of one specific TLS session. Normally, there is not a good way to get premaster keys to a Decoder in time for the parsing step. However, both Chrome and Firefox can write the premaster keys they generate to a file. This is useful for testing purposes. To configure your browser to do this, create an environment variable called `SSLKEYLOGFILE` and assign it the pathname of a file to which the keys will be written. The Decoder will accept the file exactly as written and will use all the decryption keys in the file for any encrypted traffic it captures.

This is a sample `NwConsole` script that uploads the file to a Decoder:

```
login <decoder>:50004 <username> <password>
send /decoder sslKeys --file-data=SSLKeys.txt
```

This is an example using a `curl` command (with the RESTful port) to upload the file to a Decoder:

```
curl -u "<username>:<password>" -H "Content-Type: application/octet-stream" --
data-binary @"/path/SSLKeys.txt" -X POST
"http://<hostname>:50104/decoder?msg=sslKeys"
```

After the symmetric keys are uploaded, they will immediately be used for any necessary decryption. Symmetric keys are stored in memory and there is a limit to how many can be stored at any point in time. As more keys are added, the earliest keys will be aged out. You can also add premaster keys by just passing the `random` and `premaster` parameters to `sslKeys`.

Private Keys or PEM Files

Private keys are normally stored in PEM files and are the asymmetric keys generated by services that accept TLS traffic. These keys are used during the TLS handshake to encrypt the premaster symmetric key that will be used for the rest of the payload encryption.

For example, if you have a web server whose traffic you want visibility into, you need to upload the private key it uses to encrypt traffic. You only need to do this once, as it is stored permanently (or until removed by a delete command). Private keys are automatically encrypted before storing to protect them. After upload, you must issue a parser reload command so that the newly installed key becomes visible to the HTTPS parser. Now, all TLS handshakes that use that private key will be able to be decrypted by the Decoder.

Note: Not all ciphers suites use a "known" private key (for example, Ephemeral Diffie Hellman). Encrypted traffic with those ciphers cannot be decrypted unless the premaster key is uploaded to the Decoder before the session is parsed.

These are some sample commands that upload a PEM file to be used for decryption.

Using NwConsole:

```
send /decoder sslKeys pemFilename=MyKey.pem --file-data=/path/MyKey.pem
```

Using the RESTful interface (you must provide the pemFilename parameter in the URL):

```
curl -u "<username>:<password>" -H "Content-Type: application/octet-stream" --data-binary @"/path/MyKey.pem" -X POST "http://<hostname>:50104/decoder?msg=sslKeys&pemFilename=MyKey.pem"
```


Step 6: Reload HTTPS Parser on Decoders

After uploading the private keys, you must issue a parser reload command so that the newly-installed key becomes visible to the HTTPS parser. All TLS handshakes that use that private key are now able to be decrypted by the Decoder.

To issue a parser reload command:

1. Go to  (Admin) > **Services** and select the Decoder.
2. Click  > **View** > **Explore** > **decoder**, right-click **parsers** and select **Properties**.

3. In the Properties for Decoder (DECODER) /decoder/parsers pane, select reload from the drop-down menu and click **Send**.



Step 7: View Decrypted Sessions

All storage of the packets is still encrypted, so when you search for any decrypted traffic, your query should still focus on `service=443` or `HTTPS` (HTTP over SSL/TLS). If any packets were able to be decrypted, `tls.premaster` metadata is associated with the session that has the premaster key required for decrypting that session. The session will also have crypto metadata denoting the cipher suite, with `crypto` formatted as `tls.premaster`.

Viewing Unencrypted Traffic

If packets are decrypted during the parse stage, encrypted packets are written to disk, and the matching premaster key used for decrypting is written to the `tls.premaster` meta key, analysts can view the unencrypted packets using the `tls.premaster` meta key.

One Decoder API that you can use to see the unencrypted packets is the `/sdk/content` RESTful service. You need to know the Session ID of the encrypted packets and the `flags` parameter masked to the value 128 (or 0x80 in hex). Point your browser to the Decoder RESTful interface and type in the following command, substituting the actual Session ID for `<id>`:

```
http://<decoder>:50104/sdk/content?session=<id>&flags=128&render=text
```

The Decoder returns a simple web page showing the packets after they are decrypted.

If you want to see what the packets look like encrypted, type in one of the following commands, substituting the Session ID for `<id>`:

```
http://<decoder>:50104/sdk/content&session=<id>&render=text
```

```
http://<decoder>:50104/sdk/content&session=<id>&flags&render=text
```


For more information on the `/sdk/content` service, see the manual page for `/sdk content`.

Alternatively, you can use the NetWitness Investigator thick client to decrypt and display the network sessions, as it can communicate to a Concentrator or have PCAPs and private keys loaded locally.

If you are not able to view the encrypted sessions, check [Troubleshooting](#). However, if you are able to see the `tls.premaster` metadata, that confirms that the Decoder is able to see inside the encrypted session, which enables parsers to see the unencrypted packet payload and create metadata accordingly. For these decrypted sessions, look for additional metadata around usernames, passwords, or HTTP body content that would be obfuscated when encrypted, which would result in metadata not being generated.

Troubleshooting

If you are still not able to decrypt or view the decrypted payloads after attempting the steps above, try these steps.

1. Validate that your private key is using an acceptable cipher-suite.
2. Make sure your Network Decoder is capturing the traffic matching the server for which you have the private key. To check, you can upload the private key and PCAP into Wireshark and see if it is able to decrypt it.

Decryption of Secure SMTP

In version 11.3, NetWitness Platform supports opportunistic SSL/TLS decryption, which addresses RFC 3207 (<https://tools.ietf.org/html/rfc3207>). You can add an HTTPS parser option that provides a .csv list of destination ports of the session where the STARTTLS command will be searched, with at least one encryption key that has been uploaded. This enables the STARTTLS functionality.

Note: The .csv list must include all the destination ports of a session. If a destination port is not in the list, the search for STARTTLS will not be started.

The STARTTLS search is limited to the first 1024 bytes of the session's payload, to prevent performance degradation. If STARTTLS is found, the parser uses the next client packet with payload as the start of a TLS negotiation.

To add an HTTPS parser option for secure SMTP decryption:

1. In the NetWitness Platform User Interface, go to  (Admin) > Services.
2. Select a Log Decoder service and click  > View > Explore.
3. In the left panel, select **decoder > parsers > config**.
4. In **parsers.options**, enter an HTTPS parser option in the following format that provides a .csv list of the destination ports of the session where the STARTTLS command is searched:

```
HTTPS="smtpsPorts=<port#1,port#2,port#n>"
```

The following image shows an example of an HTTPS parser option for secure SMTP decryption.

The screenshot displays the RSA NetWitness Platform interface. The left sidebar shows the navigation tree with 'Decoder (DECODER)' expanded, and 'parsers' > 'config' selected. The main panel shows the configuration for the decoder. The 'Parsers Options (parsers.options)' section is highlighted, showing the following configuration:

Configuration Item	Value
Entropy (log2=true)	GeoIP2:isp.dst,GeoIP2:isp.src,GeoIP2:latdec,GeoIP2:latdec.dst,GeoIP2:latdec.src,GeoIP2:latdec.dst
GeoIP2 (ipaddr=ip.src,ip.dst,ipv6.src,ipv6.dst)	HTTPS="smtpsPorts=25,110,465,995"

The 'HTTPS' option is highlighted with a red box, indicating the configuration for secure SMTP decryption.

Performance Considerations

Decrypting packets in real time requires extra work in the parsing stage. Before implementing this feature, plan carefully to ensure the incoming traffic bandwidth does not overwhelm the available compute power. You may need more Decoders to decrypt traffic than you would need if not decrypting.

Packets captured on a Decoder normally have a timeout of ~60 seconds in the assembly stage before they are sent to the parsing step. If the Decoder is under memory pressure due to very high bandwidth, the lifetime of the packets in Assembly may be shortened. To alleviate this situation, you can configure a longer timeout value and increase the amount of memory available to hold packets in Assembly. Also, in order to perform decryption of the packets, the Decoder must receive the decryption key before the parsing stage.

Note: Currently, only TLS 1.2 and earlier protocols can be decrypted

With no feeds loaded, the following parsers enabled, and 50% of the sessions being decrypted, a Decoder can process traffic at 3 Gbps .

Parser Name	Description
SYSTEM	Session Details
NETWORK	Network Layer
ALERTS	Alerts
GeoIP	Geographic data based on ip.src and ip.dst
GeoIP2	Geographic data by default based on IPv4 (ip.src, ip.dst) and IPv6 (ipv6.src, ipv6.dst) meta keys
HTTP	Hyper Text Transport Protocol (HTTP)
HTTP_Lua	Hyper Text Transport Protocol (HTTP) Lua
FTP	File Transfer Protocol (FTP)
TELNET	TELNET Protocol
SMTP	Simple Mail Transport Protocol (SMTP)
POP3	Post Office Protocol (POP3)
NNTP	Network News Transport Protocol (NNTP)
DNS	Domain Name Service (DNS)
HTTPS	Secure Socket Layer (SSL) Protocol
MAIL	Standard E-Mail Format (RFC822)
VCARD	Extracts VCARD fullname and email information
PGP	Identifies PGP blocks within network traffic
SMIME	Identifies SMIME blocks within network traffic
SSH	Secure Shell (SSH)
TFTP	Trivial File Transfer Protocol (TFTP)
DHCP	Dynamic Host Configuration Protocol (DHCP and BOOTP)
NETBIOS	Extracts NETBIOS computer name information.
SNMP	Simple Network Management Protocol (SNMP)
NFS	Network File System (NFS) protocol
RIP	Routing Information Protocol (RIP).
TDS	MSSQL and Sybase database protocol (TDS)
TNS	Oracle database protocol (TNS)

Parser Name	Description
IRC	Internet Relay Chat (IRC) protocol
RTP	Real Time Protocol (RTP) for audio/video
SIP	Session Initiation Protocol (SIP)
H323	H.323 Teleconferencing protocol
SCCP	Cisco Skinny Client Control Protocol
GTalk	Google Talk (GTalk)
VlanGre	Vlan ID and GRE/EtherIP tunnel addresses
BITTORRENT	BitTorrent File Sharing Protocol
FIX	Financial Information eXchange Protocol
GNUTELLA	Gnutella file sharing protocol
IMAP	Internet Message Access Protocol
MSRPC	Microsoft Remote Procedure Call protocol
RDP	Remote Desktop Protocol
SHELL	Command Shell Identification
TLSv1	TLSv1
SearchEngines	A parser that extracts search terms
FeedParser	External Feed Parser

TLS Certificate Hashing

The Network Decoder can produce hashes of certificates that are seen in the packet stream. These hashes are the SHA-1 value of any DER-encoded certificate encountered during a TLS handshake. The hashes produced can be used to compare network traffic with hashes from public SSL blacklists, such as the one from sslbl.abuse.ch.

The TLS Certificate hashing feature is disabled by default. To enable TLS Certificate hashing:

1. Go to  (Admin) > **Services**, select a Network Decoder service and  > **View** > **Explore**.
2. Select **decoder** > **parsers** > **config**.
3. In the values column next to `parsers.options`, type `HTTPS="cert.shal=true"`.
When this option is enabled, the SHA-1 is stored as a text value in the `cert.thumbprint` meta key.
4. The change takes effect on parser reload. In the left panel, right-click **parsers** and click **Properties**.
In the drop-down menu, select **reload** and then click **Send**

The SHA-256 hash of the certificate can be enabled by adding the parser option:

```
HTTPS="cert.sha256=true"
```

to a Network Decoder's `/decoder/parsers/config/parsers.options` configuration.

When this option is enabled, SHA-256 is stored as a text value in the meta keys:

```
cert.sha256  
cert.thumbprint
```

JA3 and JA3S TLS Fingerprints

The Network Decoder can produce the JA3 value of TLS clients and the JA3S value of TLS servers that are observed in a Network session. The values that are produced conform to the values generated by the open source JA3 tools (<https://github.com/salesforce/ja3>). The JA3 features are disabled by default.

To enable JA3 features:

1. In the NetWitness Platform User Interface, go to  (Admin) > **Services**.
2. Select a Decoder service and click  > **View** > **Explore**.
3. In the left panel, select **decoder** > **parsers** > **config**.
4. In **parsers.options**, add `HTTPS="ja3=true ja3s=true"`.
5. In the left panel, right click on **parsers** and select **properties**.
6. In the right panel, in the lower pane, set the value in the first drop down to **Reload**.
7. Click **Send**.

JA3 and JA3S can be enabled independently using these options. When enabled, the JA3 is stored as a text value in the meta key `ja3`. The JA3S is stored as a text value in the meta key `ja3s`.

Additional meta keys are activated using a mechanism similar to the JA3S and JA3S meta items. For example, `ja3.orig` is enabled by adding `ja3.orig=true` to the HTTPS parser options. This is also true for `ja3s.orig`, `tls.extensionlen`, and `tls.cipherlen`.

`ja3.orig`, `tls.extensionlen`, and `tls.cipherlen` are only created if JA3 is enabled.

`ja3s.orig` is only created if JA3S is enabled.

Perform Simultaneous Ingestion of the Encrypted and Decrypted Traffic Streams to Decoder

In previous versions, ingesting encrypted and decrypted traffic streams of the same session to a Decoder causes packets from both the sessions to mix on the Decoder as reassembly is performed by the same thread. It cause session parsing and content extraction inaccuracies.

In NetWitness Platform 11.6, using multi-adapter capture and multi-thread assembler features, you can configure the Decoder to receive encrypted and decrypted streams on separate adapters. The multi-thread assembler feature allows decoder to assemble packets from its corresponding capture work thread. It keeps the packets from encrypted and decrypted sessions separate during assembly and avoid inaccuracies in session parsing and content extraction.

To configure the Decoder for the simultaneous ingestion of the encrypted and decrypted traffic streams:

1. Configure Decoder to capture from multiple adapters simultaneously, see [\(Optional\) Multiple Adapter Packet Capture](#).
2. Set `/decoder/config/ assembler.threading.enabled=on` to enable multi-thread assembler.
3. Restart the Decoder service.

Community ID

The Network Decoder generates Community ID flow hash values that are compatible with the Community ID specification defined by <https://github.com/corelight/community-id-spec>.

The Community ID is written to a text-formatted meta key named `community`.

The Community ID is generated on sessions containing TCP over IPv4, TCP over IPv6, UDP over IPV4, and UDP over IPV6.

Edit Decoder System Configuration

When a service is first added to NetWitness Platform, default values for the system configuration parameters are in effect. In most cases, the default values for compression, statistics update interval, and number of threads in the thread pool are set at a good point for optimal system performance. You do not need to edit these settings unless an RSA Customer Support technician advises you to change them.


System Configuration	
Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☐
SSL Port	56004
Stat Update Interval	1000
Threads	20

One parameter that you may want to change for your environment is the SSL setting, which by default is not enabled. When enabled, the security of data transmission is managed by encrypting information and providing authentication with SSL certificates.

To edit system configuration parameters for a Decoder or Log Decoder:

1. Go to  (Admin) > Services.

2. Select a Decoder or Log Decoder service, and select  > **View > Config**.
The Services Config view for the service is displayed with the General tab open.



The screenshot displays the RSA NetWitness Platform interface for configuring a Decoder service. The top navigation bar includes links for Investigate, Respond, Users, Hosts, Files, Dashboard, and Reports. The main menu shows various service categories, with 'SERVICES' selected. The 'Config' tab is active, showing the 'General' configuration for a Decoder service.

System Configuration

Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☒
SSL Port	56004
Stat Update Interval	1000
Threads	20

Decoder Configuration

Name	Config Value
Adapter	
Berkeley Packet Filter	
Capture Interface Selected	packet_mmap_eth0 (bpf)
Cache	
Cache Directory	/var/netwitness/decoder/cache
Cache Size	4 GB
Capture Settings	
Assembler Maximum Size	32 MB
Assembler Minimum Size	0

Parsers Configuration

Specify if relevant meta data is generated to disk (Enabled), generated only in memory for other Decoder content use (Transient), or not generated at all (Disabled).

Name	Config Value
ALERTS	Enabled
DHCP	Enabled
DNS	Enabled
Entropy	Disabled
FeedParser	Enabled
FTP	Enabled
GeolIP2	(Mixed)
GTalk	Enabled
H323	Enabled
HTTP	Enabled
HTTP2	Enabled
HTTPS	Enabled
IRC	Enabled
JSNormalize	Enabled
MAIL	Enabled
NETBIOS	Enabled

Buttons: **Apply**

Footer: RSA NETWITNESS PLATFORM 11.5.0.0

3. Under **System Configuration**, click in a field that you want to edit (**Compression**, **Port**, **SSL FIPS Mode**, **SSL Port**, **Stat Update Intervals**, or **Threads**). Type a new value.
4. When finished editing, click **Apply**.
The settings become effective immediately.


Enable CPU Usage Statistics for Installed Content

The Decoder provides CPU utilization statistics for all installed content, which you can use to reveal how much CPU time is used by parsers, feeds, application rules, and lexical scanning. The statistics are visible as Stat nodes in the service tree from the Explore view when `/decoder/parsers/config/detailed.stats` is enabled and the Decoder is capturing the stats.

Each piece of content is accounted as a single percentage value (0-100) regardless of the number of parse threads running. The percentage represents an average of the CPU utilization for the content across all threads.

To enable usage statistics monitoring:

1. In the NetWitness Platform User Interface, go to  (Admin) > Services.
2. Select a Decoder service and click  > View > Explore.
The Explore view for the Decoder is displayed.
3. Open `/decoder/parsers/config/` and in the right pane, select the `detailed.stats` parameter.



Path	Value
<code>/decoder/parsers/config</code>	Decoder
Detailed Parser Statistics (detailed.stats)	no
Disabled Feeds (feeds.disabled)	
Filename Meta Generation (filename.meta)	2
Flex Enabled (Deprecated) (flex.enabled)	no
Flex Instruction Limit (flex.instruction.limit)	976.56 KB
Lua Debugger Auto Detach (lua.debugger.auto.detach)	yes
Lua Default Allocator (lua.default.allocator)	yes
Lua Enabled (lua.enabled)	yes
Lua Instruction Limit (lua.instruction.limit)	976.56 KB
Parse Maximum Bytes (parse.bytes.max)	128 KB
Parse Minimum Bytes (parse.bytes.min)	1 KB
Parse Transaction Mode (parse.transaction.mode)	off
Disabled Parsers (parsers.disabled)	WLAN,SEARCH,SMB,Entropy,GeolIP2:isp.dst,GeolIP2:isp,GeolIP2:isp.src,GeolIP2:latdec,GeolIP2:latdec.dst,Ge...
Parsers Options (parsers.options)	Entropy="log2=true" GeolIP2="ipaddr=ip.src,p.dst,ipv6.src,ipv6.dst"
Maximum Session Meta (session.meta.max)	8 KB

4. Change the value to **yes**. If the Decoder is not capturing data, start capture.
When you open the Decoder Stats node in the Explore view, the new statistic is visible.

Enable Parser Mappings

This topic tells administrators how to enable event source mapping on a Log Decoder.

The Log Collector discovers the event source type on a per-message basis. If the correct parser is not identified for the event source, a small percentage of logs may be misidentified. The misclassified messages do not populate event source rules and alerts, and the reports do not have the correct data. If there are multiple event source types associated with an IP address, it makes it difficult for the parsers to identify the exact event source from which the logs are generated.

If you map an IP address to its event source type, the Log Decoder can identify the event source from which the log is generated. When messages are delivered to the Log Decoder from a mapped event source, only the assigned parsers are queried to find event matches.

You can assign event source types to IPV4, IPV6, or the hostname value of the event source. You can also assign multiple event source types to a single IP address. You can also use the Log Collector ID when different event source types with the same IP address are sent to different Log Collectors.


Note: You can also enable parser mapping functions by navigating to  (Admin) > **Event Sources** > **Discovery**.

Enable IP Address to Event Source Mapping

To enable an IP address to event source mapping:

1. Go to  **(Admin)** > **Services** and select a Log Decoder.
2. Select  > **View** > **Config**.
3. In the Configuration page, select the **Parser Mappings** tab.

The Parser Mappings tab is displayed in the Services Config view.




Update IP to Event Source Mapping

To update an IP to event source mapping:

1. Go to  (Admin) > **Services**.
2. Select a **Log Decoder**, and in the **Actions** column, select  > **View** > **Config**.
The Services Config view is displayed.
3. Select the **Parsers Mapping** tab.
4. Click .

The Mapping Editor is displayed.



5. Any of the following mappings can be defined:

- **One Host and One Event Source Type**

In the **Host** field, enter the hostname.

For example: 10.0.0.1

- In the **Event Source(s)** field, enter the event source type.

For example: apache

- **One Host and One or More Event Source Types**

In the **Host** field, enter the hostname.

For example: 10.0.0.1

- In the **Event Source(s)** field, enter the event source type.

For example: apache, sap, aix

- **One Host, One Log Collector, and One Event Source Type**

In the **Host** field, enter the hostname and Log Collector ID.

For example: 10.0.0.1, LC-1

- In the **Event Source(s)** field, enter the event source type.

For example: apache

- **One Host, One Log Collector ID, and One or More Event Source Types**

In the **Host** field, enter the hostname and Log Collector ID.

For example: 10.0.0.1, LC-1

- In the **Event Source(s)** field, enter the event source type.

For example: apache, sap, aix

Note: The event source types are processed in the order you enter the parsers and if one or more parsers matches a log, the first parser in the list is queried. The Host/IP can be IPv4, IPv6, or Hostname.

6. Click **OK**.
The Parser Mapping is added.
7. To cancel the parser mappings selection, click **Cancel**.

Read IP to Event Source Type Mappings


To read an IP to event source type mappings:

1. Go to  (**Admin**) > **Services**, and select a Log Decoder service.
2. In the Actions column, select  > **View** > **Config**.

The Services Config view is displayed.

3. Select the **Parsers Mapping** tab.

The mappings are displayed.



Edit IP to Event Source Type Mappings

To edit IP to event source type mappings:

1. Go to  (**Admin**) > **Services**, and select a Log Decoder service.
2. In the Actions column, select  > **View** > **Config**.
The Service Config view is displayed.
3. Select the **Parser Mappings** tab.
4. Select the mapping you want to edit.

Note: You can only edit one mapping at a time.

5. Click .
6. In the **Event Source(s)** field, modify the event source(s).

Note: The host is not editable and the field is disabled.

7. Click **OK** to accept the edited Event Source.
8. To cancel the changes, click **Cancel**.

Delete IP to Event Source Type Mappings

To delete IP to event source type mappings:

1. Go to  (**Admin**) > **Services**, and select a Log Decoder service.
 2. In the Actions column, select  > **View** > **Config**.
The Service Config view is displayed.
 3. Select the **Parser Mappings** tab.
 4. Select the mapping you want to delete.
 5. Click .
- The mapping is deleted and the grid is refreshed.
6. To cancel the changes, click **Cancel**.

Sort the Hostname or Event Source Type

To sort the hostname or event source type:

1. Go to  (**Admin**) > **Services**, and select a Log Decoder service.
2. In the Actions column, select  > **View** > **Config**.
The Service Config view is displayed.
3. Select the **Parser Mappings** tab.
4. To sort a column, click in the column header.

Event Source Types are applied for your selected IP address. Logs are parsed against the parsers in the order they are listed.

Import IP to Event Source Mapping Entries


To import IP to event source mapping entries:

1. Go to  (**Admin**) > **Services**, and select a Log Decoder service.
2. In the Actions column, select  > **View** > **Config**.

The Service Config view is displayed.

3. Select the **Parser Mappings** tab.
4. Select **Actions** > **Import**.

The Import dialog is displayed.



5. Click .
6. Select the file you want to import and click **OK**.
7. To load the parser, click **Import**.

Note: You can only import one .csv file at a time.

Export IP to Event Source Mapping Entries

To export IP to event source mapping entries:


1. Go to  (**Admin**) > **Services**, and select a Log Decoder service.
2. In the Actions column, select  > **View** > **Config**.

The Service Config view is displayed.

3. Select the **Parser Mappings** tab.
4. Select the mappings you want to export.

5. Select **Actions > Export > Selection**.

The Export Selection dialog is displayed.

A screenshot of the 'Export Selection' dialog box. It has a dark gray header bar with the title 'Export Selection' and a close button (X) on the right. Below the header is a white area with a text input field. To the left of the input field is the label 'Enter File Name'. At the bottom of the dialog is a dark gray bar containing two buttons: 'Cancel' and 'Export'.

6. Enter the file name and click **Export**.

Search IP to Event Source Mapping Entries

To search IP to event source mapping entries:

1. Go to  (**Admin**) > **Services**, and select a Log Decoder service.
2. In the Actions column, select  > **View > Config**.

The Service Config view is displayed.

3. Select the **Parser Mappings** tab.
4. In the Parsers Mappings toolbar, enter the Host or Event Source in the **Filter** field.
5. Click **Enter**.

The Hosts or Event Sources that match the names entered in the **Filter** field are displayed.

Enable or Disable Lua and Flex Parsing Systems

This topic tells administrators how to enable or disable Lua and Flex parsing systems on a Decoder or Log Decoder. Flex parsers are deprecated and disabled by default.

The settings to enable or disable Lua and Flex parsing systems are configured correctly by default and you do not typically have to change them. However, you may need to adjust these settings at the request of RSA Customer Care or for troubleshooting purposes.

In addition to configuring individual parsers, you can enable and disable all Lua parsing as well as all Flex parsing in the Services Explore view. You enable and disable the Lua parsing and Flex parsing systems settings separately, but they work in the same way.

- If you **disable** the Lua or Flex parsing system, the corresponding parsing system is disabled and no parsers are loaded.
- If you **enable** the Lua or Flex parsing system, the corresponding parsing system is enabled and individual parsers are enabled and disabled following the current individual configurations.

To enable or disable Lua and Flex parsing systems on a Decoder or Log Decoder:

1. Go to  (Admin) > Services.
2. Select a Decoder or Log Decoder and  > View > Explore.
The Services Explore view for the selected service is displayed.

3. In the Node list, navigate to and select `/decoder/parsers/config`.

Path	Value
Detailed Parser Statistics (detailed.stats)	no
Disabled Feeds (feeds.disabled)	
Filename Meta Generation (filename.meta)	2
Flex Enabled (Deprecated) (flex.enabled)	no
Flex Instruction Limit (flex.instruction.limit)	976.56 KB
Lua Debugger Auto Detach (lua.debugger.auto.detach)	yes
Lua Default Allocator (lua.default.allocator)	yes
Lua Enabled (lua.enabled)	yes
Lua Instruction Limit (lua.instruction.limit)	976.56 KB
Parse Maximum Bytes (parse.bytes.max)	128 KB
Parse Minimum Bytes (parse.bytes.min)	1 KB
Parse Transaction Mode (parse.transaction.mode)	off
Disabled Parsers (parsers.disabled)	WLAN,SEARCH,SMB,Entropy,GeolIP2:isp.dst,GeolIP2:isp,GeolIP2:isp.src,GeolIP2:latdec,GeolIP2:latdec.dst,Ge...
Parsers Options (parsers.options)	Entropy="log2=true" GeolIP2="!paddr=ip.src,ip.dst,ipv6.src,ipv6.dst"
Maximum Session Meta (session.meta.max)	8 KB

4. In the values panel:
- To enable the Lua parsing system, in the value field for `lua.enabled`, type **yes**.
 - To disable the Lua parsing system, in the value field for `lua.enabled`, type **no**.
 - To enable the Flex parsing system, in the value field for `flex.enabled`, type **yes**.
 - To disable the Flex parsing system, in the value field for `flex.enabled`, type **no**.

Map IP Address to Service Type for Log Parsing

This topic describes the procedure to map an IP address to a service type for log parsing.

The Log Collector discovers event source type on a per-message basis. If the correct parser is not used for the specific event source, the messages that are common between event source types are misclassified. The misidentified messages will not populate service rules and alerts, and the reports will not have proper information. Also, if there are multiple services associated with an IP address, it can be difficult for the parsers to identify the exact service from which the log is generated.

If you map an IP address to its services, the Log Decoder can identify the service from which the log is generated. When messages come into the log decoder from a mapped service, the assigned parsers are loaded to find event matches.

You can assign service types to IPV4, IPV6 or hostname value of the event source. You can also assign multiple service types to a single IP address. You can also use the CollectorID when different service types with the same IP address are sent to different collectors.

Map an IP Address to a Service Type

To map an IP address to a service type, do the following:

1. Go to  (Admin) > Services.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select  > **View** > **Explore**.
3. Go to `/decoder/parsers` node, right-click **parsers**, and select **Properties**.
4. In the **Properties** view, specify the **ipdevice** command with the following parameters:
`op=add/remove entries="ipaddress=service"`
 for example, `op=add entries="10.100.201.300=ciscoasa"`
5. Click **Send**.

Properties for Log Decoder (Log Decoder) /decoder/parsers.	
ipdevice ▼	Parameters <code>op=add entries=""</code> <code>=rhlinux</code> <code>=ciscoasa,rhlinux"</code>
Message Help	
Map IP to Device type in log parsing. Multiple device types mapped to the same ip/host are prioritized in the order in which they are listed. Takes effect immediately. security.roles: parsers.manage parameters: op - <string, {enum-one:add edit remove describe}> The operation to perform.	
Response Output	
IP2Device entry edited	

IPdevice Command

In the `ipdevice` command, three operations are available:

- add: This operation adds or updates entries in the ipdevice map. Multiple space delimited address/type pairs may be specified.

op=add entries="**address=service type**"

- remove: This operation removes entries from the ipdevice map. Multiple space delimited address/type pairs may be specified.

op=remove entries="**address**"

- describe: This operation returns the values currently in the ipdevice map.

Map an IP Address to a Time Zone

This topic discusses RSA NetWitness Platform support for Event Time.

Often times logs do not fully specify timestamps and may be missing time zone information. To properly normalize such timestamps to UTC, the Log Decoder provides the ability to associate devices from a specific address (IPv4 or IPv6) or hostname to a time zone or a fixed offset.

Timestamp consideration for log files:

- Some logs are in UTC format.
- Some logs that are not in UTC include a timezone offset value to adjust the time accordingly.
- For logs that use local time, with no offset provided, you can create a source mapping to manually adjust times for logs from that event source.

Three time zone formats are currently accepted and are shown in the following examples:

- Olson format: `America/Anguilla`
- POSIX format: `AST2:45ADT0:45,M4.1.6/1:45,M10.5.6/2:45`
- Offset by Minutes format: `= -500`

NetWitness Platform maps the device address (IPv4 or IPv6) or hostname to a specific time zone or offset. Event time meta that is parsed from a log that is from a mapped address and does not include an offset or time zone as part of the timestamp is adjusted to UTC according to the mapping.

To map an IP address to a time zone, do the following:

1. Go to  (Admin) > **Services**.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select  > **View** > **Explore**.
3. Go to `/decoder/parsers` node, right click **Parsers**, and select **Properties**.
4. In the **Properties** view, specify the `iptmzone` command with the following parameters:
`op=add entries="ipaddress=timezone"`
for example: `op=add entries="10.10.10.10=Africa/Addis Ababa"`
5. Click **Send**.

tzinfo Command

You can view the strings that can be used to specify the time zone by using the `tzinfo` command.

To view the time zone strings, do the following:

1. Go to  (Admin) > **Services**.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select  > **View** > **Explore**.
3. Go to `/decoder/parsers` node, right click **Parsers**, and select **Properties**.

4. In the **Properties** view, specify the `tzinfo` command with the following parameters:

```
op=tznames
```

5. Click **Send**.

The list of time zone strings is returned in the **Response Output** text box.

iptmzone Command

In the `iptmzone` command, three operations are available:

- **add**: This operation adds or updates entries in the `iptmzone` map. Multiple space delimited address/type pairs may be specified.

```
op=add entries="address=time zone"
```
- **remove**: This operation removes entries in the `iptmzone` map. Multiple space delimited address/type pairs may be specified.

```
op=remove entries="address"
```
- **describe**: This operation returns the values currently in the `iptmzone` map.

Examples

The following examples provide instances for mapping IP addresses to time zones:

- If you want to map two different entries with different IPV4 values and time zone, enter the following parameter in the **iptmzone** command and click **Send**

```
op=add entries="10.10.10.10=America/Anguilla 10.10.10.11=Pacific/Rarotonga"
```
- If you want to remove an entry for a single IPV4 value and time zone, enter the following parameter in the **iptmzone** command and click **Send**.

```
op=remove entries="10.5.245.9"
```
- If you want to create a single entry for an IPV6 value and time zone, enter the following parameter in the **iptmzone** command and click **Send**.

```
op=add entries="2001:DB8:85A3::8A2E:370:7334=America/Anguilla"
```
- If you want to create a single entry to map an IPV4, IPV6, or hostname with the Minute Offset, Olson, or POSIX format, enter the following string in the **iptmzone** command and click **Send**.

```
op=add entries="10.168.0.2=America/Anguilla  
2001:DB8:85A3::8A2E:370:7334=0500 nwappliance21=EST5EDT,M3.2.0/2,M11.1.0"
```

Change the Date format

The Log Decoder parsing engine uses the `event.time` meta key to keep track of the time an event occurs for incoming log messages. You can use the `mdformat` for the **iptmzone** command to change the date format. The Log Decoder currently has the ability to change the dates in the logs to have the following format:

```
mdy or dmy (month/day/year or day/month/year)
```

Notes:

- Event time metadata in a parser does not account for `dmy` and `mdy`.
- Currently, there is no functionality to detect this scenario and adjust parsing automatically.

To change the date format, do the following:

1. Go to  (Admin) > **Services**.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select;  > **View** > **Explore**.
3. Go to `/decoder/parsers` node, right click **Parsers**, and select **Properties**.
4. In the **Properties** view, select **iptmzone** from the drop-down list, and specify the command with the following parameters:

`op=add entries="<ipaddress>" mdformat=[dmy | mdy | none]`

for example: `op=add entries="1.1.1.1" mdformat=dmy`

Where:

- `ipaddress` is the Device IP address, and
- `mdformat` can be **dmy**, **mdy**, or **none**.

Properties for  - Log Decoder (Log Decoder) /decoder/parsers.

iptmzone	Parameters	op=add entries=1.1.1.1 mdformat=dmy

Message Help

Map IP to time zone in log parsing. Takes effect after parser reload.

security.roles: parsers.manage

parameters:

op - <string, {enum-one:add|edit|remove|describe}> The operation to perform (edit|describe).

Response Output

IP2TimeZone entry edited

5. Click **Send**.

NetWitness Platform maps the IP address along with the date format in the Log Decoder. Event time meta items are updated according to their respective mappings.

Event Time Filter Example

If you want to see logs that have an event time between 3:02 PM and 3:15 PM on May 17, 2019, create the following filter:

Edit Filter event.time > '2019-May-17 15:02:00' && event.time < '2019-May-17

☐ Simple
 ☒ Advanced
 ☐ Recent

event.time > '2019-May-17 15:02:00' && event.time < '2019-May-17 15:15:00'

Cancel OK Reset

Note: The query builder may tell you “invalid expression,” but that is a validation error in the User Interface. The query works because even though this is a text (string) search, the system actually compares the ASCII values. So, for example the expression **A < B** would be true.

Below is a sample of the event meta values returned from the query:



Missing Year Support

If there is no year associated with a format string, the parser attempts to populate the year heuristically. In most cases, the value is populated based on the current year. The current year is assigned as the message year (UTC), as per the clock on the Log Decoder.

The following are the various other scenarios:

1. Latent log during transition to new year.
2. Logs from forward time zones (or skewed clocks) just before new year transition.
3. Latent logs received where a leap day cannot be successfully assigned to an appropriate leap year.

Latent log during transition to new year

If the day is more than 31 days into the future, the year is decremented.

For example,

The message date is Dec-31. The current date is 2018-Jan-1. The temporary assigned date, 2018-Dec-31, will be more than 31 days into the future.

This will cause the year component to be decremented to 2017, resulting in a reasonable message time.

Logs from forward time zones (or skewed clocks) just before new year transition

If the day is more than 334 days into the past, the year is incremented.

For example,

The message date is Jan-1. The current date is 2017-Dec-31. The temporary assigned date, 2017-Jan-1, will be more than 334 days into the past.

This will cause the year component to be incremented to 2018, resulting in a reasonable message time.

Latent logs received where a leap day cannot be successfully assigned to an appropriate leap year

If the day (Feb 29) is invalid (say, the assigned year is NOT a leap year), the relative position to the current time cannot be calculated. To do this, the year is decremented in an attempt get a valid time stamp.

Because the position of the leap day relative to the transition to the new year, along with the expectation that no logs would be received this far into the future, we do not ever make an attempt to increment the year to produce a valid time stamp.

After the year is decremented, the same logic is then followed (refer to statement 1 and 2). If the result is still an invalid day. The parser throws and a valid message time cannot be parsed.

Limitations

Note: Currently, there is no mechanism to assign a year to the import of logs.

As this pertains to leap days, if you find the correct leap year, an inconsistent data would result. For example, if a set of messages is two years old and during a leap year, only a single day would remain which is accurate. Because of the heuristic described above, the remaining set of messages would have time stamps one year too.

For data consistency, usually you cannot find valid leap years beyond 1 month and less than 11 months.

Examples

The following examples provide instances for logs with year and logs without year:

- If the log is with year, event time is generated as below:

```
%emcavamar: 1704^^2017-04-07^^07:50:02^^1^^<event-source NodeID="avamar"
ProgramName="com.avamar"/>^^1270626^^SYSTEM^^ERROR^^OK^^MCS:DPN_
Proxy^^Internal server error^
<MESSAGE
    id1="1:01"
    id2="1"
    eventcategory="1605020000"
    functions="&lt;@msg:*PARMVAL ($MSG) &gt;&lt;@event_time:*EVNTTIME
($MSG, '%W-%G-%F %H:%U:%O', fld2, fld3) &gt;"
    content="&lt;fld1&gt;^^&lt;fld2&gt;^^&lt;fld3&gt;^^&lt;fld4&gt;^^
&lt;&lt;event-source NodeID=&quot;&lt;hostname&gt;&quot;
ProgramName=&quot;&lt;fld8&gt;&quot;/&gt;^^&lt;fld9&gt;^^&lt;cate
gory&gt;^^&lt;severity&gt;^^&lt;event_
type&gt;^^&lt;agent&gt;^^&lt;event_description&gt;" />
```

- If the log is without year, event time is still generated, The current year is assigned as the message year (UTC), as per the clock on the Log Decoder.

```
%emcavamar: 1704^^04-07^^07:50:02^^1^^<event-source NodeID="avamar"
ProgramName="com.avamar"/>^^1270626^^SYSTEM^^ERROR^^OK^^MCS:DPN_
Proxy^^Internal server error^
<MESSAGE
    id1="1:01"
    id2="1"
    eventcategory="1605020000"
    functions="&lt;@msg:*PARMVAL ($MSG) &gt;&lt;@event_time:*EVNTTIME
($MSG, '%G-%F %H:%U:%O', fld2, fld3) &gt;"
    content="&lt;fld1&gt;^^&lt;fld2&gt;^^&lt;fld3&gt;^^&lt;fld4&gt;^^
&lt;&lt;event-source NodeID=&quot;&lt;hostname&gt;&quot;
ProgramName=&quot;&lt;fld8&gt;&quot;/&gt;^^&lt;fld9&gt;^^&lt;cate
gory&gt;^^&lt;severity&gt;^^&lt;event_
type&gt;^^&lt;agent&gt;^^&lt;event_description&gt;" />
```

Obtain Log Files from a Pre-11.0 Log Decoder

NetWitness 11.0. added the capability to view a small sampling of recent logs for specific devices through detail tabs of the Discovery View. By default, Log Decoders prior to 11.0 do not have the necessary configuration to enable this feature, but a few minor changes can make it available.

To enable logs preview for a pre-11.0 Log Decoder, follow these steps on the Log Decoder:

1. Go to  (Admin) > **Services** > select a **Log Decoder**, then select  > **View** > **Config**.
2. Click the **Files** tab and select **index-logdecoder-custom.xml** from the drop-down menu.
3. Add the following three lines at the end of the file (before the closing language tag):


```
<key description="Device IP" level="IndexValues" name="device.ip" format="IPv4" valueMax="100000"
defaultAction="Open"/>
<key description="Device IPv6" level="IndexValues" name="device.ipv6" format="IPv6"
valueMax="100000" defaultAction="Open"/>
<key description="Device Host" level="IndexValues" name="device.host" format="Text"
valueMax="100000" defaultAction="Open"/>
```

4. Click **Apply**.
5. Restart the Log Decoder service as follows.

Select Log Decoder service > **Explore** > **decoder** > **Properties** > **reset**. You select **reset** from a drop down menu.

Click **Send** after you select reset.

This is an example of the **index-logdecoder-custom.xml** file.



Note: Discovery Scores are only available for 11.x and above Log Decoders. Discovery Scores for pre-11.x Log Decoders are displayed as Unavailable.

The following example shows the Discovery Score as **Unavailable** in the **Details** view for a pre-11.0 Log Decoder.

The screenshot shows the RSA Event Sources interface. The top navigation bar includes tabs for Hosts, Services, Event Sources (selected), Endpoint Sources, Health & Wellness, System, and Security. Below this, there are sub-tabs for Discovery, Manage, Monitoring Policies, Alarms, and Settings. The main content area is divided into two sections: Filters on the left and Event Sources on the right.

Filters:

- Event Source: Contains 32.226.1
- Event Source Type: Please select an event source type
- Show Acknowledged: ☐
- Mapping Type:
 - ☒ All
 - ☒ None
 - ☒ Auto
 - ☒ Manual

Event Sources:

Event Source	Discovery Score	Acknowl	Mapping Type	Log Collector(s)	Log Decoder(s)	Event Source Type(s)
32.226.1.36	Unavailable	No	None		LogDecoder - 3522f8a0416c469c96e0b879af4ad664	oracle rhlinux solaris
32.226.1.2	Unavailable	No	None		LogDecoder - 3522f8a0416c469c96e0b879af4ad664	msexchange msias winevent_nic
32.226.1.1	Unavailable	No	None		LogDecoder - 3522f8a0416c469c96e0b879af4ad664	msexchange mssql symantecav winevent_n
32.226.1.3	Unavailable	No	None		LogDecoder - 3522f8a0416c469c96e0b879af4ad664	checkpointfw1

Page 1 of 1 | Page Size 50

Note: Device logs are only available for 11.x and above Log Decoders.

The following example shows the message that is displayed in the Logs panel for a pre-11.0 Log Decoder.

The screenshot shows the RSA Event Source Type(s) for '12.22.23.12'. The top navigation bar includes tabs for Hosts, Services, Event Sources (selected), Endpoint Sources, Health & Wellness, System, and Security. Below this, there are sub-tabs for Discovery, Manage, Monitoring Policies, Alarms, and Settings. The main content area is divided into two sections: Event Source Type(s) for '12.22.23.12' on the left and Logs on the right.

Event Source Type(s) for '12.22.23.12':

Event Source Type	Discovery Score
bigfix	Unavailable

Logs:

Timestamp	Log Decoder	Discovery Score	Message
-	3522f8a0416c469c96e0b879af4ad664	-	Discovery logs view is only available for 11.x and above Log Decoders by default. See documentation (link?) for enabling on earlier versions.

Attributes:

Log Collector	Log Decoder
3522f8a0416c469c96e0b879af4ad664	3522f8a0416c469c96e0b879af4ad664
UPS Protected	false

Upload a Log File to a Log Decoder


There are occasions when you want to analyze a log file that is not available on the service you are using. You can upload a log file captured on another service to NetWitness Platform. Log filenames are of the type **.log**.

When a log file is uploaded to a Log Decoder, the Log Decoder analyzes and generates meta for each log it contains. These logs are added to the already decoded logs on the Log Decoder and are available for analysis. NetWitness Platform includes a filename tracking option that makes searching for a particular set of logs easier. When the log file is uploaded with file tracking, the Log Decoder adds meta to each log based on the uploaded filename. You can then filter sessions for analysis using that meta.

The option to upload a log file is dimmed when other Log Decoder operations prevent an upload from occurring, for example, when the Log Decoder is capturing logs.

To import a log file to a Log Decoder:

1. Go to  (Admin) **Services**.
2. Select a Log Decoder in the **Service** grid, and select  > **View** > **System**.
The Services System view for the Log Decoder is displayed.
3. In the toolbar, click **Upload Log File**.



4. To choose a log file, click **Browse**.
A directory view is displayed.
5. Select the log file that you want to upload.
The filename is displayed in the **Upload File** field.
6. If you want the Log Decoder to add meta to the logs based on the filename, click the checkbox next to **Track Filename**.
7. To upload the file, click **Upload**.
The selected file is uploaded and a status message indicates that the file is uploaded. The log file is available for analysis.

Upload a Packet Capture File


There are occasions when you want to analyze a packet capture file that is not available on the service you are using. You can upload a file captured on another service to NetWitness Platform. Supported packet capture file types are `pcap` and `pcap.gz`.

When a packet capture file is uploaded to a Decoder, the Decoder creates sessions from the packet capture file packets. These sessions are added to the already decoded sessions on the Decoder and are available for analysis. NetWitness Platform includes a filename tracking option that makes searching for a particular set of sessions easier. When the packet capture file is uploaded with file tracking, the Decoder adds meta to the sessions based on the uploaded filename. You can then filter sessions for analysis using that meta.

The option to upload a packet capture file is dimmed when other Decoder operations prevent an upload from occurring; for example, when the Decoder is capturing packets.

To select and upload a packet capture file:

1. Go to  **(Admin)** > **Services**.
The Administration Services view is displayed.
2. Select the Decoder name, and  > **View** > **System**.
The Services System view for the Decoder is displayed.
3. In the toolbar, click **Upload Packet Capture File**.
The **Upload Packet Capture File** dialog is displayed.



4. To choose a capture file, click **Select**.
A directory view is displayed.
5. Browse the directory and select the packet capture file that you want to upload.
The filename is displayed in the **Upload File (pcap, pcap.gz)** field.
6. If you want the Decoder to add meta to the sessions based on the filename, click the checkbox next to **Track Filename**.
7. To upload the file, click **Upload**.
A progress bar shows upload progress.

Upload time varies depending on the size of the file. When the file upload is complete, a status message is displayed. The file is now available for investigation.

Simultaneous Import and Capture

You can import or upload PCAPs while the capture is running using RESTful API. Make a note of the following points before you start importing during capture:

- The source file metadata is not created for imported PCAPs when the capture is running. If capture is stopped, the source file metadata is created as in older versions.
- The NWD files can be imported, but not while the capture is running. If you attempt to import an NWD file while the capture is running, the Decoder returns an error. An active import will block starting/stopping of import and capture.
- The file tracking meta is not created during the simultaneous import and capture.

Decoder and Log Decoder References

This is a collection of references, which provide information about the user interface for Decoders and Log Decoders in NetWitness Platform, with references to the procedures that describe the work you can do in that part of the user interface.

Topics

- [Services Config View - Capture Policies Tab](#)
- [Services Config View - Edit Capture Policies Wizard](#)
- [Services Config View - Data Privacy Tab](#)
- [Services Config View - Data Retention Scheduler](#)
- [Services Config View - Feeds Tab](#)
- [Services Config View - Files Tab](#)
- [Services Config View - General Tab](#)
- [Services Config View - Parsers Tab](#)
- [Services Config View - Parser Mappings Tab](#)
- [Services Config View - Rules Tabs](#)
- [Services System View - Decoders](#)

Services Config View - Capture Policies Tab

In the Capture Policies tab ( (Configure) > **Capture Policies**, administrators can configure selective network data collection policies.

Selective network data collection gives administrators the ability to apply centrally managed capture policies across their Network Decoders. This results in better use of Decoder resources, including hard drive space, which leads to more predictable costs and lessens the burden of managing multiple services. You can determine which traffic is stored and how it is stored by using policies. Each policy contains a list of supported base protocols and definitions for handling any other protocols that are detected. The administrator can choose to deploy predefined policies that capture:


- All base and other protocols (Full Capture - All Protocols)
- Only metadata on all base and other protocols (Capture Meta Only - All Protocols)
- Only metadata on all base protocols and drop all other protocols (Capture Meta on Base Protocols, Drop all other protocols)
- All base protocols and only metadata on all other protocols (Full Capture on Base Protocols, Meta only on all other protocols)

The predefined policies are not configurable. The only way you can edit predefined policies is by assigning services (such as Decoders) to deploy them.

Administrators can create custom policies to give further control over the deployment. A base set of protocols is available for alterations by the administrator, allowing you to choose what level of capture you prefer on a per-protocol basis. If you are only making slight changes, a good start for customization is to clone one of the predefined policies and alter it. These centrally managed policies are then applied to services (Network Decoders) to allow handling multiple use cases across your environment.

Workflow

The following figure depicts the workflow for common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	use predefined selective collection policies*	(Optional) Configure Selective Network Data Collection
Administrator	create new policies from predefined ones*	(Optional) Configure Selective Network Data Collection

User Role	I want to...	Documentation
Administrator	create custom policies*	(Optional) Configure Selective Network Data Collection
Administrator	verify policies	(Optional) Configure Selective Network Data Collection
Administrator	unpublish policies*	(Optional) Configure Selective Network Data Collection
Administrator	delete policies*	(Optional) Configure Selective Network Data Collection

*You can complete these tasks here.

Related Topics

- [Services Config View - Edit Capture Policies Wizard](#)

Quick Look

This is an example of the Capture Policies tab.

- 1 Create New - Provides a wizard to create a new policy.
- 2 Edit - Provides a wizard to edit an existing custom policy.
- 3 Clone - Selects an existing policy (predefined or custom) and copies it to create a new policy.
- 4 Delete - Removes policies.

The following table describes the contents of the Capture Policies tab.

Column	Description
Policy Name	The name of the policy (not editable for predefined policies).
Policy Description	The description of the policy (not editable for predefined policies).

Column	Description
Publication Status	• Unpublished: policies that have not been deployed to Decoders yet. • Unpublished Edits: policies that have been updated but are not deployed to Decoders. Clicking Save and Close after making edits keeps them from automatically being published, allowing the administrator to wait until a specific outage window to publish them. This would display as Unpublished (for policies not previously published to Decoders) or Unpublished Edits (for policies that have been previously published to Decoders, but the updates have not been deployed to Decoders). • Published: policies which have been deployed to Decoders. • Failed: policies that failed to execute, for example, if a Decoder is offline or a system is down. For information about how to troubleshoot failed policies, see "Troubleshooting Policy Deployments That Fail" in (Optional) Configure Selective Network Data Collection.
Service Assignment	Identity of the Decoder or service on which the policy is run.
Policy Updated	Date the policy was edited.
Updated By	Role of the person who updated the policy.
Policy Created	Date on which the policy was created.
Created By	Role of the person who created the policy.

Services Config View - Edit Capture Policies Wizard


In the Edit Policies view ( (Configure) > **Capture Policies** > **Edit**, administrators can configure selective network data collection policies. The panels for creating a new policy and editing an existing policy are the same except that in the editing wizard, the definitions of the existing policy are displayed.

You can edit existing policies to customize them for your environment. A base set of protocols are available for alterations, allowing you to choose what level of capture you prefer on a per-protocol basis. If you are only making slight changes, a good start for customization is to clone one of the predefined policies and alter it.

Note: A set of predefined policies are available immediately. The only way you can edit a predefined policy is by assigning Decoders to the policy to collect the targeted data.

Workflow

The following figure shows the workflow for common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	create new policies from predefined ones*	(Optional) Configure Selective Network Data Collection
Administrator	create custom policies*	(Optional) Configure Selective Network Data Collection
Administrator	unpublish policies*	(Optional) Configure Selective Network Data Collection
Administrator	delete policies*	(Optional) Configure Selective Network Data Collection

*You can complete these tasks here.

Related Topics

- [Services Config View - Capture Policies Tab](#)

Quick Look

1. To access the Edit Policy wizard, go to  (Configure) > **Capture Policies**, select a policy, and click **Edit**.
The Identify Policy page is displayed, where you can change the name and description of the policy.
2. To view the Define Policy page, click **Next**.

In the Define Policy page, you can define rule actions for the protocols for which you are interested in collecting data. The rule actions are:

- **Collect Meta Only:** Collect metadata
- **Drop All:** Drop metadata and network packets
- **Collect All:** Collect metadata and network packets

You can click **Save and Close** to save your policy without deploying it on Decoders. The following table describes the actions of the buttons on this page.

Button	Description
Previous	Go to the previous page of the wizard
Next	Keep updates from the current page and go to the next page of the wizard.
Save and Close	Save the policy definitions without publishing the policy.
Save and Publish	Save the policy definitions and publish the policy to Decoders to begin data collection.
Cancel	Cancel the definitions and return to the Capture Policies tab.

3. Click Next to go to the Assign Policy page.

This is where you assign the policy to Decoders to capture the data. The following table describes the columns in this page.

Column	Description
Service Name	The name of a service. For 11.5, Decoder is the only available service.
Policy	The policy that a Decoder is assigned to.
Host	The IP address or name of the host system for the Decoder.
Service Type	The type of service the policy is published to. For 11.5, the service type is decoder .

Column	Description
Version	The version of NetWitness Platform that is installed on the service host.

Services Config View - Data Privacy Tab

In the Data Privacy tab ( (Admin) > **Services** > **Select a Decoder or Log Decoder** >  > **Config** > **Data Privacy tab**), administrators can configure data privacy parameters for certain Core services. For the Decoder and Log Decoder, you can set the default hash algorithm and salt.


What do you want to do?

User Role	I want to ...	Documentation
Administrator	configure hash algorithm and salt	"Configure the Hash Algorithm and Salt" in the <i>Data Privacy Management Guide</i> . (Go to the Master Table of Contents to find all RSA NetWitness Platform 11.x documents.)

Related Topics

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)

Quick Look



RSA Investigate Respond Users Hosts Files Dashboard Reports

HOSTS SERVICES EVENT SOURCES ENDPOINT SOURCES HEALTH & WELLNESS SYSTEM SECURITY

Change Service | Decoder | Config

General Files Data Retention Scheduler Network Rules App Rules Correlation Rules Feeds Parsers **Data Privacy**

Configure Hash Algorithm and Salt

Modify default behavior for hashing of meta key values at the time of creation by defining a default hash algorithm and/or salt.

Hash Algorithm

The algorithm to use for any keys with a transform defined that does not specify type. Use Duplicate in validating that expected hashing behavior is occurring in the network.

Salt

The optional salt value prepended to any value that is hashed (sha-256). Best practice is to configure a salt, which ensures that obfuscated values are unique to this site and therefore more difficult to deduce.

Apply

RSA NETWITNESS PLATFORM 11.5.0.0

The Data Privacy tab has the Configure Hash Algorithm and Salt configuration settings. The following table describes the parameters in this tab.

Parameter	Description
Hash Algorithm	Displays a drop-down list of hash algorithms to use for any keys with a transform that does not specify algorithm type. Possible values are SHA-256 and Duplicate. Duplicate is a special algorithm available for administrators to use when validating that expected hashing behavior is occurring in the network.
Salt	Indicates the optional salt value prepended to any value that is hashed. Best practices for security purposes dictate a salt value that is no less than 100 bits or 16 characters in length. Configuring a value ensures that obfuscated values are unique to this site and therefore more difficult to deduce. For more information on this field, see "Configure Data Obfuscation" in the <i>Data Privacy Management</i> guide.
Apply	Applies any changes.

Services Config View - Data Retention Scheduler

In the Data Retention Scheduler tab, you can set the rollover criteria for removing database records from primary storage using an age-based threshold. You can also schedule the timing to check whether the threshold is reached.

To access the Data Retention Scheduler tab, go to  (Admin) > **Services** > select a **Decoder** or **Log Decoder** service and click   > **View** > **Config** > **Data Retention** tab.

What do you want to do?

User Role	I want to...	Documentation
Administrator	Schedule the timing to see if the threshold is reached.	Configure Transaction Handling on a Decoder

Related Topics

- [Configure Common Settings on a Decoder](#)
- [Decoder and Log Decoder Quick Setup](#)

Quick Look

This is an example of the Data Retention Scheduler tab.

Define the rollover criteria for removing database records from primary storage using an age-based threshold, and schedule the timing for checking if the threshold has been reached.

Threshold **1** → ☒ Duration ☐ Date **2**

Days Hours Minutes

Run **3** → ☒ Interval ☐ Date & Time **4**

Hours Minutes

Apply **Reset**

RSA NETWITNESS PLATFORM 11.5.0.0

- 1 Threshold Duration:** Removes database files older than the selected number of days, minutes, or hours.
- 2 Threshold Date:** Removes database files older than the selected UTC date (YYYY-MM-DD-HH:MM:SS) that are not compatible with minutes, hours, or days parameters.
- 3 Run Interval:** Indicates the number of hours between executions.
- 4 Run Date and Time:** Defines which days of the week to execute the scheduler, as well as time of execution in HH:MM:SS format for the local time of the service.

Services Config View - Feeds Tab

Feeds and parsers are Lua programs loaded and compiled when either processing capture files in NetWitness Investigate or capturing data with Decoders. Most commonly, they are used for static meta extraction and service identification.

Note: Pre-11.0 versions of NetWitness used FLEXPARSE programs in addition to Lua programs; Flexparsers are deprecated in NetWitness Platform 11.0. Unless otherwise stated, any reference to Decoders applies to Log Decoders as well.

NetWitness Platform uses feeds to create metadata based on externally defined meta values. A feed is a list of data that is compared to sessions as they are captured or processed. For each match, additional metadata is created. This data can identify and classify malicious IPs or incorporate additional information such as department and location based on internal network assignments. Some examples of feeds include threat feeds to identify BOTNets, DHCP mappings, or even active directory information such as physical location or logical department.

Feeds can be added, removed, and updated while a Decoder is running without affecting capture. The

Feeds tab ( (Admin) > **Services** > select a Decoder or Log Decoder service and click   > **View** > **Config** > **Feeds tab**) provides a user interface for managing feeds on Decoders.

What do you want to do?


User Role	I want to...	Documentation
Administrator	configure feeds	Configure Parsers and Feeds
Administrator	enable and disable parsers	Enable and Disable Parsers and Log Parsers

Related Topics

- [Configure Common Settings on a Decoder](#)
- [Decoder and Log Decoder Quick Setup](#)
- [Upload Feeds Dialog](#)

Quick Look

This is an example of the Feeds tab.



- 1 Feeds Tab Toolbar - Provides options to work with feeds in the grid
- 2 Feed List - Lists all feeds that are currently deployed on the Decoder

Feeds Tab Toolbar

Feature	Description
 Upload	Displays the Upload Feeds dialog.
	Deletes the selected feeds.

Feeds List

The Feeds list provides a listing of all currently deployed feeds for the Decoder.

Column	Description
Name	The name of the feed or the feed file.

Column	Description
Live	Indicates if the feed originated from Live. Possible values are Yes , No , or N/A . • Yes = Installed through Live • No = Installed through NetWitness Platform • N/A = The feed has no attributes file created by NetWitness Platform to track the installation date. The feed may have been installed manually, not through NetWitness Platform or Live Services. Manually installed feeds still function properly.
Date Installed	The date the feed was pushed to the service.

Upload Feeds Dialog

This topic describes the features of the Upload Feeds dialog in the Services Config view > Feeds tab.

The **Upload** option in the Services Config view > Feeds tab displays the Upload Feeds Dialog, in which you can manage the uploading of feeds to a Decoder or Log Decoder.

What do you want to do?


User Role	I want to...	Documentation
Administrator	prepare a list of feeds for upload	Edit, Upload, or Remove a Feed
Administrator	view and delete upload jobs	Edit, Upload, or Remove a Feed

Related Topics

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)

Quick Look

This is an example of the Upload Feeds dialog.



File List

The File list is the place to prepare a list of feeds for uploading. You can add files from a directory structure, and delete files from the list if you decide that you don't want to upload a particular file. When the list is ready, clicking **Upload** starts the upload process.

Feature	Description
	Opens a view of the directory structure where you can select files to add to the File list.
	Deletes the selected files from the File list.
File Name	Lists the feed files you have added from a file system in preparation for uploading to a Decoder. When you click Upload , the files listed here are uploaded.

Upload Job List

The Upload Job list provides a view of upload jobs started by clicking **Upload**.

Feature/Column	Description
 Delete	Deletes an upload job.
Progress	Displays progress of an upload job.
Start Time	Displays the start time of an upload job.
File Name	Lists filename of the feed being uploaded.
Status	Displays the status of upload job.

Upload Feeds Dialog Buttons


Feature	Description
Cancel	Closes the Upload Feed dialog.
Upload	Starts uploading the feed files listed in the File list. Each feed is listed in a separate row in the Upload job list.

Services Config View - Files Tab

The Decoder and Log Decoder configuration files are visible and editable in the  (Admin) > Services > Config view > Files tab. "Edit Core Services Configuration Files" in the *Hosts and Services Getting Started Guide* provides general instructions for editing files. (Go to the [Master Table of Contents](#) to find all RSA NetWitness Platform 11.x documents.)

Like other Core services, both the Decoder and Log Decoder have an index file, and may also have a crashreporter, netwitness, and scheduler. The Decoder and Log Decoder index files are named **index-decoder-custom.xml** and **index-logdecoder-custom.xml**.

Note: Table-map.xml and table-map-custom.xml are available only for Log Decoders with log content installed.



What do you want to do?

User Role	I want to...	Documentation
Administrator	obtain log files from pre-11.0 Log Decoder	Obtain Log Files from a Pre-11.0 Log Decoder

User Role	I want to...	Documentation
Administrator	edit files and parsers	• Feed Definitions File • Flex Parsers • GeoIP2 Parsers • Lua Parsers • Decoder Snort Detection • Search Parser • Wireless LAN Configuration

Related Topics

- [Configure Common Settings on a Decoder](#)
- [Decoder and Log Decoder Quick Setup](#)
- [Create Custom Meta Keys Using a Custom Feed](#)

Quick Look


Filename	Description
GeoPrivate.ipl	This fixed parser takes the IP addresses and converts them to geographical locations. The locations are displayed through the Google Earth display.
feed-definitions.xml	Used to create custom feeds, this is the XML schema used by the Decoder to define a feed message when it creates a .feed file.
traffic_flow_options.lua	Used to provide directionality information. Update this file with environment-specific internal and external subnets for the Lua parser to create proper directionality in metadata. The parser is described in RSA Content for RSA NetWitness Platform .
search.ini	This is the Search Parser configuration file. The Search Parser is a custom parser, used to generate metadata by scanning for pre-defined keywords and regular expressions.
wlan-config.xml	This is the wireless LAN configuration file (9/9/2009). This file controls the 802.11 parsers. Its chief purpose is to control decryption of raw 802.11 frames captured by the Decoder.

Services Config View - General Tab

The General tab for a Decoder in the Services Config view provides a way to manage basic service configuration, configure data capture, and select the parsers that are applied to the captured data. To access the General tab, go to **Admin > Services** > select a Decoder or Log Decoder and click  > **View > Config > General tab**.

Workflow

The following figure depicts common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	configure capture settings*	Configure Capture Settings
Administrator	manage parsers and log parsers*	Enable and Disable Parsers and Log Parsers
Administrator	start and stop data capture	Start and Stop Data Capture
Administrator	configure rules	Configure Decoder Rules

*You can complete these tasks here.

Related Topics

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)
- [Configure Parsers and Feeds](#)

Quick Look

The first figure is an example of the General tab for a Decoder. The second is the General tab for a Log Decoder.

System Configuration

Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☒
SSL Port	56004
Stat Update Interval	1000
Threads	20

Decoder Configuration

Name	Config Value
Adapter	
Berkeley Packet Filter	
Capture Interface Selected	packet_mmap_eth0 (bpf)
Cache	
Cache Directory	/var/netwitness/decoder/cache
Cache Size	4 GB
Capture Settings	
Assembler Maximum Size	32 MB

Parsers Configuration

Specify if relevant meta data is generated to disk (Enabled), generated only in memory for other Decoder content use (Transient), or not generated at all (Disabled).

Name	Config Value
ALERTS	Enabled
apt_artifacts	Enabled
Avamar	Enabled
BGP_lua	Enabled
bittorrent_lua	Enabled
Canon_BJNP	Enabled
cerber	Enabled
china_chopper	Enabled
creditcard_detection_lua	Enabled
CustomTCP	Enabled
db2_lua	Enabled
DCERPC	Enabled
Derusbi_Server_Handshake	Enabled
DHCP	Enabled
DHCP_lua	Enabled
DNP3_lua	Enabled

Apply

System Configuration

Name	Config Value
Compression	0
Port	50002
SSL FIPS Mode	☐
SSL Port	56002
Stat Update Interval	1000
Threads	20

Log Decoder Configuration

Name	Config Value
Adapter	
Berkeley Packet Filter	
Capture Interface Selected	log_events,Log Events
Cache	
Cache Directory	/var/netwitness/logdecoder/cache
Cache Size	4 GB
Capture Settings	
Assembler Maximum Size	32 MB

Parsers Configuration

Specify if relevant meta data is generated to disk (Enabled), generated only in memory for other Decoder content use (Transient), or not generated at all (Disabled).

Name	Config Value
ALERTS	Enabled
apt_artifacts	Enabled
Avamar	Enabled
BGP_lua	Enabled
bittorrent_lua	Enabled

Service Parsers Configuration

Specify if relevant meta data is generated to disk (Enabled), generated only in memory for other Decoder content use (Transient), or not generated at all (Disabled).

Name	Config Value
accurev	☒
actionvantage	☒
activity	☒
aforecloudlink	☒
airdefense	☒
airmagnet	☐
airtightmc	☒
alix	☒

Apply

1 System Configuration - Manages service configuration for a Decoder or a Log Decoder.

- 2** Decoder Configuration or Log Decoder Configuration - Lets you view and edit service configuration parameters for a Decoder or Log Decoder.
- 3** Parsers Configuration - Lets you select parsers to use on the Decoder or Log Decoder.
- 4** Service Parsers Configuration (Log Decoders only) - Lets you select service parsers to use on the Log Decoder.

System Configuration Section

The System Configuration section manages service configuration for a Decoder. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change.

System Configuration	
Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☐
SSL Port	56004
Stat Update Interval	1000
Threads	20

The System Configuration section has these parameters.

Parameter	Description
Compression	The minimum number of bytes that must be transmitted per response before compression. A setting of 0 disables compression. The default value is 0 . A change in value is effective immediately for all subsequent connections.
Port	Determines the port used by the service. Note: If you change the port number, ensure that you restart the service.
SSL FIPS mode	If enabled, all the data transferred in the network will be encrypted using SSL.
SSL Port	Indicates the port used for encrypting using SSL.
Stat Update Interval	The number of milliseconds between statistic updates on the system. Lower numbers cause more frequent updates and can slow down other processes. The default value is 1000 . A change in value is effective immediately.
Threads	The number of threads in the thread pool to handle incoming requests. A setting of 0 lets the system decide. A change takes effect on service restart.

Decoder Configuration Section

The Decoder Configuration section provides a way to view and edit service configuration parameters for a Decoder or Log Decoder. When a service is first added, default values are in effect. You can edit these values to manage traffic capture.

Decoder Configuration	
Name	Config Value
Adapter	
Berkeley Packet Filter	
Capture Interface Selected	
Cache	
Cache Directory	/var/netwitness/decoder/cache
Cache Size	4 GB
Capture Settings	
Assembler Maximum Size	32 MB
Assembler Minimum Size	0
Assembler Session Flush	1
Assembler Session Pool	50000
Assembler Timeout Packets	60

Scrolling to the bottom of the section reveals these additional Decoder Configuration parameters.

Decoder Configuration	
Name	Config Value
Assembler Timeout Session	60
Capture Autostart	☐
Capture Buffer Size	32 MB
Parse Maximum Bytes	128 KB
Parse Minimum Bytes	1 KB
Parse Threads	0
 Database Max File Sizes	
Meta File Size	3 GB
Packet File Size	4 GB
Session File Size	512 MB
 Hash	
Hash Directory	

Adapter Section

Adapter parameters configure the network interface for capture as described in [Configure Capture Settings](#).

Cache Section

Cache parameters configure the cache directory and size for session cache files. The following table describes the cache settings. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change.

Cache Parameter	Description
Cache Directory	The directory where session cache files are stored. The default value is <code>/var/netwitness/decoder/cache</code> . Change takes effect immediately.
Cache Size	The maximum size, in Megabytes (MB), that all files in the cache directory can attain before the oldest files are deleted. Once the threshold is reached, the cache size is reduced by 10%. The default value is 4 GB . Change takes effect immediately.

Capture Settings Section

The Capture Settings section provides a way to configure operational capture settings. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change.

Capture Settings Parameter	Description
Assembler Maximum Size	Specifies the maximum size in bytes that a session's packet data size can attain. The default value is 32 MB . Change takes effect immediately.
Assembler Minimum Size	Specifies the minimum size in bytes that a session must have in order to generate metadata. A value of 0 means every session has metadata generated. The default value is 0 . Change takes effect immediately.
Assembler Session Flush	Specifies whether a session is removed from the assembler when the session's last chain is removed from the assembler. The default value is 1. • 2 = if the first packet of a session times out of assembler, the session is removed from assembler after parsing is complete. Any subsequent packets for this session create a new session in assembler. • 1 = If the last chain of a session times out of assembler, the session is removed from assembler. Any subsequent packets for this session create a new session in assembler. • 0 = If the last chain of a session times out of assembler, the session is left in assembler until it times out. Any subsequent packets for this session are filtered. Change takes effect on service restart.
Assembles Session Pool	Specifies the number of entries in the session pool. The default value is 350000 . Change takes effect on service restart.
Assembler Timeout Packets	Specifies the number of seconds before a packet or chain is timed out. The default value is 60 . Change takes effect immediately.
Assembler Timeout Session	Specifies the number of seconds before a session is timed out. Default value is 60 . Change takes effect immediately.
Capture Autostart	Specifies whether capture begins automatically each time Decoder is started. When checked, the value = yes. When unchecked, the value = no. The default value is no . Change takes effect immediately.
Capture Buffer Size	The capture memory buffer allocation in Megabytes. Default value is 64 MB . Change takes effect on service restart.
Parse Maximum Bytes	The maximum number of bytes to scan a stream for additional tokens. When the first token is found, the stream is scanned up to the set number of bytes, but no further. A setting of 0 removes the early termination and the full stream is scanned regardless of size. The default value is 128 KB . Change takes effect immediately.
Parse Minimum Bytes	The minimum number of bytes to scan a stream for the first token. If no token is found within the set number of bytes, scanning is terminated. A setting of 0 removes the early termination and the full stream is scanned regardless of size. The default value is 1 KB . Change takes effect immediately.

Capture Settings Parameter	Description
Parse Threads	The number of parse threads to use for session parsing. A value of 0 means let the server decide. The default value is 0 . Change takes effect on service restart.

Database Max File Sizes Section

The Database Max File Sizes section controls the maximum file size for various databases. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change.

File Size Parameter	Description
Meta File Size	The maximum size of meta database files in Megabytes. The default value is 10 MB . Change takes effect on service restart.
Packet File Size	The maximum size of packet database files in Megabytes. The default value is 10 MB . Change takes effect on service restart.
Session File Size	The maximum size of session database files in Megabytes. The default value is 100 MB . Change takes effect on service restart.

Hash Section

The Hash section settings control data base file hashing options. There is a small performance penalty when hashing.

Hash Parameter	Description
Hash Directory	The server directory where all hash files are written. If empty, each hash file is written to the same directory as the file being hashed. The default value is blank. Change takes effect on service restart.

Parsers Configuration Panel

The Parsers Configuration panel provides a way to select parsers to use on the Decoder. Within some parsers, you can also configure the metadata that the parser creates. See [Enable and Disable Parsers and Log Parsers](#) for detailed information and procedures.

Parsers Configuration		Enable All	Disable All
Specify if relevant meta data is generated to disk (Enabled), generated only in memory for other Decoder content use (Transient), or not generated at all (Disabled).			
Name	Config Value		
☒ ALERTS	Enabled		
☒ DHCP	Enabled		
☒ DNS	Enabled		
☒ Entropy	Disabled		
FeedParser	Enabled		
☒ FTP	Enabled		
☒ GeoIP2	(Mixed)		
☒ GTalk	Enabled		
☒ H323	Enabled		
☒ HTTP	Enabled		
☒ HTTPS	Enabled		
☒ IRC	Enabled		
☒ MAIL	Enabled		
☒ NETBIOS	Enabled		
☒ NETWORK	Enabled		
NFS	Enabled		
☒ NNTP	Enabled		
☒ PGP	Enabled		
☒ POP3	Enabled		

Service Parsers Configuration Section for Log Decoder

The Service Parsers Configuration section provides a way to select Service parsers to use on the Log Decoder.

Service Parsers Configuration		Enable All	Disable All
Name	Config Value		
accurev	☒		
actiancevantage	☒		
actividentity	☒		
aforecloudlink	☒		
airdefense	☒		
airmagnet	☐		
airtightmc	☒		
aix	☒		
alcatelomniswitch	☒		
apache	☒		
apachetomcat	☐		

Services Config View - Parsers Tab

In the Services Config View > Parsers tab, you can view deployed parsers on a Decoder or Log Decoder, upload parsers, and delete deployed parsers. Parsers can be added and removed while a Decoder or Log Decoder is running without affecting capture.

To access the Parsers tab, go to  (Admin) > **Services** > select a **Decoder** or **Log Decoder** service and click   > **View** > **Config** > **Parsers** tab.

What do you want to do?

User Role	I want to...	Documentation
Administrator	View deployed parsers.	Enable and Disable Parsers and Log Parsers
Administrator	Upload parsers to a Decoder or Log Decoder.	Enable and Disable Parsers and Log Parsers

Related Topics

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)
- [Upload and Delete Custom Parsers](#)

Quick Look

This is an example of the Parsers tab. The Parsers grid lists all parsers that are currently deployed on the Decoder.

- 1 Name:** The name of the parser or the parser file.
- 2 Live:** Indicates if the parser originated from Live. Possible values are **Yes**, **No**, or **N/A**.
 - **Yes** = Installed through Live Services.
 - **No** = Installed through NetWitness.
 - **N/A** = The parser has no attributes file created by NetWitness to track the installation date. The parser may have been installed manually, not through NetWitness or Live Services.
- 3 Date Installed:** The date the parser was pushed to the service.

Parsers Tab Toolbar

The Parsers Tab toolbar has options to work with parsers in the grid.

Feature	Description
 Upload	Enables you to upload parsers to a Decoder or Log Decoder.
	Requests confirmation that you want to delete the selected parsers. You can select No to cancel the deletion, or select Yes to delete the selected parsers.

Services Config View - Parser Mappings Tab

This topic provides a description of the configurable options for a Log Decoder in the Parser Mappings tab.

In the Parser Mappings Administrators can configure log parser mappings for Log Decoder services. To access the Parser Mappings tab, go to  (Admin) > **Services** > select a Log Decoder service and click  > **View** > **Config** > **Parser Mappings** tab.

Note: You can also configure log parser mappings for Log Decoder services by navigating to **Admin** > **Services** > **Event Sources** > **Discovery**.

This feature is intended to track a subset of event sources that is parsing against the wrong parser.

What do you want to do?

User Role	I want to...	Documentation
Administrator	Manage IPs for Event Source Mapping.	Enable Parser Mappings

Related Topics

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)

Quick Look

This is an example of the tab.

This feature allows you to map log messages from individual event sources to specific log parsers. This can be useful when logs are being parsed against an incorrect parser.

1 Parser Mappings Toolbar - Provides options to work with parser mappings in the list

2 Parser Mappings List - Lists all parsers that are currently mapped on the Log Decoder

Parser Mappings Toolbar

The Parser Mappings toolbar has options to work with parser mappings in the grid.

Feature	Description
	Add a parser mapping.
	Delete the selected parser mapping.
	Edit a parser mapping.
	Refresh the list of parser mappings.
	Display the Actions menu. Import - Import a parser mapping to a file. Export - Save a parser mapping to a file.

Parser Mappings List

The Parser Mappings list displays all parsers that are currently mapped on the Log Decoder.

Parameter	Description
Host	Displays the IP address of the host.
Event Source	Displays the event sources that are parsing incorrectly.

Parser Mappings Editor Dialog

The Parser Mappings Editor dialog allows you to update an IP to event source mapping.

To access the Parser Mappings Editor dialog, in the Services Config view for a Log Decoder, select the Parser Mappings tab.

Data Export Tab

The Data Export tab (**Admin > Services >** select a Decoder or Log Decoder and click  > **View > Config > Data Export** tab) exports meta data and raw logs from the Decoder or Log Decoder and converts it to Logstash JSON object by configuring the export connector through Logstash.

Prerequisites

Make sure you have log collector service to configure Logstash in the UI.

What do you want to do?


User Role	I want to...	Documentation
Administrator	Configure Export Connector	Configure Logstash Event Sources in NetWitness Platform

Related Topics

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)
- [Configure Decoder Rules](#)
- [Services Config View - Rules Tabs](#)

Quick Look

The following figure shows an Data Export tab.



The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. The main navigation bar has tabs for 'HOSTS', 'SERVICES', 'EVENT SOURCES', 'ENDPOINT SOURCES', 'HEALTH & WELLNESS', 'SYSTEM', and 'SECURITY'. The 'SERVICES' tab is active, showing a list of services including 'endpointloghybrid1 - Log Decoder'. The 'Config' button is visible. Below the service list, there are tabs for 'General', 'Files', 'Data Retention Scheduler', 'App Rules', 'Correlation Rules', 'Feeds', 'Parsers', 'Parser Mappings', 'Data Privacy', 'Data Export', and 'Appliance Service Configuration'. The 'Data Export' tab is selected. The main content area displays a table with columns 'Host' and 'Name'. The table lists four Log Collectors: 'endpointloghybrid1 - Log Collector', 'logcollector - Log Collector', 'loghybrid1 - Log Collector', and 'VLC - Log Collector'. The bottom of the page shows a footer with 'RSA NETWITNESS PLATFORM' and version '11.6.0.0'.

The data export list displays all the Log Collectors in your deployment.


Parameter	Description
Host	Displays the IP address of the host.
Name	Displays the name of the Log Collector.

Services Config View - Rules Tabs

The Rules tabs in the Services Config view ( (Admin) > **Services** > select a service and click  > **View** > **Config**) enable you to define and manage capture rules. Each type of rule has a list with slightly different columns and different parameters in the Rule Editor dialog. Application and correlation rules apply to both Decoders and Log Decoders. Network rules apply only to Network Decoders.

Workflow

The following figure depicts the workflow for common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	configure capture settings	Configure Capture Settings
Administrator	manage parsers and log parsers	Enable and Disable Parsers and Log Parsers
Administrator	start and stop data capture	Start and Stop Data Capture
Administrator	configure rules*	Configure Decoder Rules
Administrator	import, export, or push a rule*	Configure Decoder Rules
Administrator	enable or disable a rule*	Configure Decoder Rules
Administrator	add, edit, or delete a rule*	Configure Decoder Rules


*You can complete these tasks here.

Related Topics

- [Configure Common Settings on a Decoder](#)
- [Decoder and Log Decoder Quick Setup](#)
- [App Rules Tab](#)
- [Correlation Rules Tab](#)
- [Network Rules Tab](#)

Quick Look

This is an example of the App Rules tab.




- 1 Rules Tab Toolbar - Provides options to work with rules in the list
- 2 Rules Actions Menu - Provide options to manage sets of rules
- 3 Rules List Context Actions - Displays the Rules List Context Menu

Rules Tab Toolbar

The toolbar is the same for all Config view > Rules tabs.




Feature	Description
Actions	Displays the Actions menu.
	Adds a new rule to a service.
	Deletes a rule from a service.
	Allows rule modification.
	Disables a rule (without deleting the rule).
	Enables (reactivates) a rule.

Feature	Description
Filter	The input field for a search string. NetWitness Platform filters the rules dynamically as you type a search string. Clicking x clears the input field, restoring the unfiltered view.
Apply	Saves the changes made to rules and applies the configured rules to a service. Until you apply changes, it is possible to reload the rules as they were before current modifications.
Revert	Discards unsaved changes to the list and reverts to the unedited rules.

Rules Actions Menu

The Actions menu has options that help to manage sets of rules.



Option	Description
Import	Imports a set of rules into the user interface so that it can be applied to a service. You can edit the rules before applying.
Export	Saves selected rules or all rules to an .nwr file on the client machine.
Push	Allows rules to be applied to other services (Decoders or Log Decoders) or Decoders belonging to a service group. When pushing, the rules can either be merged (update existing rules and append new ones) or replaced. • Push > All. Pushes all rules to other services. All rules on the target services are removed and replaced with all of the rules on the source service. • Push > Selection. Pushes selected rules to other services. You have two options: • Replace. Deletes all rules on the target services and replaces them with the selected rules from the source service. • Merge. Merges the selected rules with the existing rules on the target services
History	Displays the last ten snapshots of rules applied through NetWitness Platform. You can select and apply (restore) a snapshot to the Decoder at anytime.

Rules List Context Actions

Within a rules list, right-clicking a row displays the Rules list context menu.

Option	Description
Cut	Deletes the current rule.
Copy	Copies the current rule.
Paste Above	Pastes the copied rule above the current rule.
Paste Below	Pastes the copied rule below the current rule.
Edit	Edits the current rule.
Insert Below	Inserts imported rules below the current rule.
Insert Above	Inserts imported rules above the current rule.
Export Selection	Exports the selected rules.
Push Selected Rules	Pushes the selected rules to other services.

App Rules Tab

The App Rules tab (**Admin > Services > select a Decoder or Log Decoder and click  > View > Config > App Rules tab**) enables you to manage application rules. NetWitness Platform applies application rules at the session level.

What do you want to do?


User Role	I want to...	Documentation
Administrator	add or edit application rules	Configure Application Rules

Related Topics

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)
- [Configure Decoder Rules](#)
- [Services Config View - Rules Tabs](#)

Quick Look

The following figure shows an App Rules tab and the table describes the columns.



The screenshot displays the RSA NetWitness Platform interface, specifically the 'App Rules' tab under the 'Config' section for a selected decoder. The interface includes a navigation bar at the top with tabs like 'General', 'Files', 'Data Retention Scheduler', 'Network Rules', 'App Rules', 'Correlation Rules', 'Feeds', 'Parsers', 'Data Privacy', and 'Appliance Service Configuration'. Below the navigation bar, there is a table of application rules. The table has columns for 'Status', 'Name', 'Condition', 'Session Data', 'Alert', and 'Last Updated By'. The rules are listed in a table, and there are buttons for 'Apply' and 'Revert' at the bottom.

Status	Name	Condition	Session Data	Alert	Last Updated By
☐	nw110025	ip.proto=1 && (filename exists filetype exists)		alert.id	
☐	nw110030	service=0 && (filename exists filetype exists)		alert.id	
☐	nw60005	service = 53 && udp.dstport = I-52,54-u && stre...		alert.id	
☐	nw60010	service = 23 && tcp.dstport = I-22,24-u && strea...		alert.id	
☐	nw60015	service = 21 && tcp.dstport = I-20,22-u && strea...		alert.id	
☐	nw60020	service = 80 && tcp.dstport = I-79,81-u && strea...		alert.id	
☐	nw60025	service = 22 && tcp.dstport = I-21,23-u && strea...		alert.id	
☐	nw60030	service = 25 && tcp.dstport = I-24,26-u && strea...		alert.id	
☐	nw60035	service = 67 && tcp.dstport = I-66,68-u && strea...		alert.id	
☐	nw60040	service = 69 && tcp.dstport = I-68,70-u && strea...		alert.id	
☐	nw60045	service = 110 && tcp.dstport = I-109,111-u && s...		alert.id	
☐	nw60050	service = 119 && tcp.dstport = I-118,120-u && s...		alert.id	
☐	nw60055	service = 135 && tcp.dstport = I-134,136-u && s...		alert.id	
☐	nw60060	service = 137 && tcp.dstport = I-136,138-u && s...		alert.id	
☐	nw60065	service = 139 && tcp.dstport = I-138,140-444,44...		alert.id	
☐	nw60070	service = 161 && tcp.dstport = I-160,162-u && s...		alert.id	
☐	nw60075	service = 443 && tcp.dstport I= 443,993,995 && ...		alert.id	
☐	nw60080	service = 520 && tcp.dstport = I-519,521-u && s...		alert.id	
☐	nw60085	service = 1433 && tcp.dstport = I-1432,1434-u && ...		alert.id	

Column	Description
Pending	This column indicates whether a rule has pending changes. Rules that are currently active on the Decoder have no indicator. If the rule is new or has been modified, the column contains  . Once the rules are applied, the pending indicator is removed.
Name	This is the rule name, a descriptive identifier for the rule.
Condition	This is the definition of the condition that triggers an action when matched.
Session Data	This column displays the Session Data action taken when a packet matches the rule. Possible values are Filter , Keep , or Truncate .
Alert	This column displays the name of the custom alert that the Decoder generates when metadata matches the rule.
Status	This column indicates whether the rule is enabled or disabled with a circle icon. If the circle is filled green, the rule is enabled. If the circle is empty, the rule is disabled.

Rule Editor Dialog

The following figure shows the Rule Editor dialog for an application rule.

Rule Editor

Rule Definition

Rule Name

Condition

*All string literals and time stamps must be quoted.
Do not quote number values and ip addresses.
Examples : 1. device.group='Windows Compliance' && service = 443
2. time = '2015-jan-01 00:00:00' - u
3. ip.src = 10.0.0.0/8,172.16.0.0/12,192.168.0.0/16 || extension = 'torrent'*

Session Data

☒ Stop Rule Processing

☐ Keep

☐ Filter

☒ Truncate

☒ All

☐ After First Bytes

☐ After SSL/TLS Handshake

NOTE: If applied to a session that is not SSL/TLS, this option will truncate the payload.

Session Options

☐ Alert ☐ Forward ☐ Transient

Alert On ▼

Reset Cancel OK

The Rule Editor dialog provides the fields and options needed to define an application rule.

Field	Description
Rule Name	The descriptive name that identifies the rule.
Condition	The definition of the condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta from the Intellisense window actions. As you build the rule definition, Intellisense displays syntax errors and warnings. All string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details.

The following table describes the Session Data actions and options.

Action	Description
Stop Rule Processing	If checked, further rule evaluation ends if the rule is matched, and the session is saved in accordance with the session action. If not checked, rule evaluation continues until all rules are evaluated.
Keep	The packet payload and associated metadata are saved when they match the rule.
Filter	The packet is not saved when it matches the rule.
Truncate	Truncate All – truncates all session payload bytes. The packet payload is not saved when it matches the rule, but packet headers and associated metadata are retained. This is the default truncation option. Truncate After First <n> Bytes – truncates the session payload bytes after the specified first <n> bytes, where <n> is an integer. The packet payload is not saved after <n> bytes when it matches the rule, but packet headers and associated metadata are retained. Truncate SSL/TLS After Handshake – truncates the payload for all sessions except in the case of an SSL/TLS session, where the SSL exchange is preserved, but the rest of the payload is not saved. This option is for use with SSL parsers.
Alert and Alert On	If Alert is checked, the packet generates a custom alert when metadata matches the rule. You can select the name of the alert in the Alert On field.
Forward	Enables the performance of syslog forwarding when the log matches the rule.
Transient	Prevents the alert metadata that is created from being written to the disk.

The following table describes Rule Editor dialog actions.

Action	Description
Reset	Resets the contents of the dialog to their values before editing; changes are discarded.
Cancel	Cancels any edits and closes the Rule Editor dialog.
OK	Saves the new rule or edited rule, and adds it to the rules grid. The Rule Editor dialog closes.
Save	(Rules with deprecated syntax only) Applies a corrected rule individually to the Decoder service. See Fix Rules with Invalid Syntax .

Correlation Rules Tab

The Correlation Rules tab ( **(Admin)** > select a service and click   > **View** > **Config** > **Correlation Rules tab**) enables you to manage correlation rules. Basic correlation rules are applied at the session level and alert the user to specific activities that may be occurring in their environment. NetWitness Platform applies correlation rules over a configurable sliding time window.

What do you want to do?


User Role	I want to...	Documentation
Administrator	add or edit a correlation rule	Configure Correlation Rules

Related Topics


- [Configure Common Settings on a Decoder](#)
- [Decoder and Log Decoder Quick Setup](#)
- [Configure Decoder Rules](#)
- [Services Config View - Rules Tabs](#)

Quick Look

The following figure shows the Correlation Rules tab.



The following figure shows the Rule Editor dialog for a correlation rule.



The following table describes the Correlation Rules tab columns.

Column	Description
Pending	This column indicates whether a rule has pending changes. Rules that are currently active on the Decoder have no indicator. If the rule is new or has been modified, the column contains  . Once the rules are applied, the pending indicator is removed.
Name	This is the descriptive name for the rule.
Condition	This is the definition of the condition that triggers an action when matched. In conditions, all string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details.
Instance Key	This is the target indicator to base the event upon. It can be a single primary key, such as <code>ip.src</code> or a compound primary key such as <code>ip.src, ip.dst</code> .
Threshold	This is the minimum number of occurrences required to trigger a correlation session and can include a associated key that identifies the meta type that were are counting to determine if the condition is satisfied. The correlation engine cannot use IPv4 or IPv6 as an associated meta type. Use one of these three arguments: <code>u_count(associated_key)</code> = the count of unique values of the specified key. A key is required. <code>sum(associated_key)</code> = the values of the specified key. a key is required. <code>count()</code> = number of sessions, no associated key used. If included, it is ignored.
Time Window	This is the duration in hours, minutes, or seconds within which the threshold must be reached to trigger a correlation session.
Status	This column indicates whether the rule is enabled or disabled with a circle icon. If the circle is filled green, the rule is enabled. If the circle is empty, the rule is disabled.

The **Rule Editor** dialog provides the fields and options needed to define a network rule. The fields correspond exactly to the grid columns.

Action	Description
Reset	Resets the contents of the dialog to their values before editing; changes are discarded.
Cancel	Cancels any edits and closes the Rule Editor Dialog.
OK	Saves the new rule or edited rule, and adds it to the rules grid. The Rule Editor Dialog closes.
Save	(Rules with deprecated syntax only) Applies a corrected rule individually to the Decoder service. See Fix Rules with Invalid Syntax .

Network Rules Tab

The Network Rules tab ( (Admin) > Services > select a Decoder and click  > View > Config > **Network Rules tab**) enables you to manage network rules. NetWitness Platform applies network rules at the packet level. Network rules consist of rule sets from Layer 2, Layer 3, and Layer 4. Multiple rules can be applied to the Decoder. Rules can be applied to multiple layers (for example, when a network rule filters out specific ports for a specific IP address). Network rules apply only to Network Decoders.

What do you want to do?

User Role	I want to...	Documentation
Administrator	add, edit, or fix network rules	Configure Network Rules

Related Topics

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)
- [Configure Decoder Rules](#)
- [Services Config View - Rules Tabs](#)

Quick Look

The following figure shows the Network Rules tab.

RSA

InvestigateRespondUsersHostsFilesDashboardReports

HOSTSSERVICESEVENT SOURCESENDPOINT SOURCESHEALTH & WELLNESSSYSTEMSECURITY

Change ServiceNH - DecoderConfig

GeneralFilesData Retention SchedulerNetwork RulesApp RulesCorrelation RulesFeedsParsersData PrivacyAppliance Service Configuration

+ - [icon] Enable Disable Actions

Filter

	Status	Pending	Name	Condition	Packet Data	Alert
☐	☒		upgrade1	tcp.srcport=443	keep	
☐	☒		upgrade2	ip.src=127.0.0.1	truncate	

ApplyRevert

RSA NETWITNESS PLATFORM11.5.0.0

The following figure shows the Rule Editor dialog for a network rule.

Rule Editor

Rule Definition

Rule Name: Subnet Filter

Condition: ip.addr=192.168.2.0/24

*All string literals and time stamps must be quoted.
Do not quote number values and ip addresses.
Examples : 1. ip.src = 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16
2. tcp.srcport= 20,21,22,80*

Session Data

☒ Stop Rule Processing

☐ Keep

☒ Filter

☐ Truncate

Session Options

☐ Assemble

☐ Application Meta

☐ Network Meta

☐ Alert

Reset Cancel OK

The following table describes the columns in the Network Rules grid.

Column	Description
Pending	This column indicates whether a rule has pending changes. Rules that are currently active on the Decoder have no indicator. If the rule is new or has been modified, the column contains . Once the rules are applied, the pending indicator is removed.
Name	This is the rule name, a descriptive identifier for the rule.
Condition	This is the definition of the condition that triggers an action when matched.
Packet Data	This column displays the Session Data action taken when a packet matches the rule. Possible values are Filter , Keep , or Truncate .
Alert	This column indicates whether the Decoder generates a custom alert when metadata matches the rule. Possible values are Enabled or Disabled .
Status	This column indicates whether the rule is enabled or disabled with a circle icon. If the circle is filled green, the rule is enabled. If the circle is empty, the rule is disabled.

The **Rule Editor** dialog provides the fields and options needed to define a network rule.

The following table describes the Rule Definition fields.

Field	Description
Rule Name	The descriptive name that identifies the rule.
Condition	The definition of the condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta from the Intellisense window actions. As you build the rule definition, Intellisense displays syntax errors and warnings. In conditions, all string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details. This section also describes the meta keys that NetWitness Platform supports for use in network rule conditions.

The following table describes the Session Data actions.

Action	Description
Stop Rule Processing	If checked, further rule evaluation ends if the rule is matched, and the session is saved as indicated. If not checked, rule evaluation continues until all rules are evaluated.
Keep	The packet payload and associated meta are saved when they match the rule.
Filter	The packet is not saved when it matches the rule.
Truncate	The packet payload is not saved when it matches the rule, but packet headers and associated meta are retained.

The following table describes the session options.

Action	Description
Assemble	If checked, the assembler assembles the packet chain when it matches the rule.
Network Meta	The packet generates network metadata when it matches the rule.
Application Meta	The packet generates application metadata when it matches the rule.
Alert	The packet generates a custom alert when metadata matches the rule.

The following table describes Rule Editor dialog actions.

Action	Description
Reset	Resets the contents of the dialog to their values before editing; changes are discarded.
Cancel	Cancels any edits and closes the Rule Editor dialog.
OK	Saves the new rule or edited rule, and adds it to the rules grid. The Rule Editor dialog closes.
Save	(Rules with deprecated syntax only) Applies a corrected rule individually to the Decoder service. See Fix Rules with Invalid Syntax .

Services System View - Decoders

A Log Decoder is a special type of Decoder, and is configured and managed in a similar way to a Decoder. Therefore, most of the information in this section refers to both types of Decoders. Differences for Log Decoders are noted.

To reach the Services System view, go to  (Admin) > Services > select a Decoder or Log Decoder >  > **View** > **System**.

Workflow

The following figure depicts the workflow for common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	configure capture settings	Configure Capture Settings
Administrator	manage parsers and log parsers	Enable and Disable Parsers and Log Parsers
Administrator	start and stop data capture*	Start and Stop Data Capture
Administrator	upload packet capture and log files*	Upload a Log File to a Log Decoder Upload a Packet Capture File
Administrator	reset log stats, perform host tasks, shutdown the service, shutdown the appliance service, and reboot the host*	<i>Hosts and Services Getting Started Guide</i>
Administrator	configure rules	Configure Decoder Rules

*You can complete these tasks here.

Related Topics

Go to the [Master Table of Contents](#) to find all RSA NetWitness Platform 11.x documents.

- [Decoder and Log Decoder Quick Setup](#)
- [Configure Common Settings on a Decoder](#)
- "Services System View" in the *Hosts and Services Getting Started Guide*

Quick Look

This is an example of the Services System view for a Decoder.

The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes tabs for Investigate, Respond, Users, Hosts, Files, Dashboard, and Reports. The main navigation bar has tabs for HOSTS, SERVICES (selected), EVENT SOURCES, ENDPOINT SOURCES, HEALTH & WELLNESS, SYSTEM, and SECURITY. Below the main navigation bar, there are sub-tabs: Change Service, NH - Decoder (selected), and System. A red box highlights a row of controls: Upload Packet Capture File, Stop Capture, Host Tasks, Shutdown Service, Shutdown Appliance Service, and Reboot.

Decoder Service Information

Name	NH (Decoder)
Version	11.5.0.0 (Rev null)
Memory Usage	466 MB (1.45% of 32174 MB)
CPU	2%
Running Since	2020-Aug-20 18:14:34
Uptime	3 hours 25 minutes 58 seconds
Current Time	2020-Aug-20 21:40:32

Appliance Service Information

Name	NH (Host)
Version	11.5.0.0 (Rev null)
Memory Usage	31200 KB (0.09% of 32174 MB)
CPU	2%
Running Since	2020-Aug-20 18:14:29
Uptime	3 hours 26 minutes 3 seconds
Current Time	2020-Aug-20 21:40:32

Decoder User Information

Name	admin
Groups	Administrators
Roles	aggregate, connections.manage, database.manage, decoder.manage, dpo.manage, index.manage, logs.manage, parsers.manage, rules.manage, sdk.content, sdk.manage, sdk.meta, sdk.packets, services.manage, storedproc.execute, storedproc.manage, sys.manage, users.manage

Host User Information

Name	admin
Groups	Administrators
Roles	appliance.manage, connections.manage, logs.manage, services.manage, storedproc.execute, storedproc.manage, sys.manage, users.manage

Session Information

Session	User	IP Address	Login Time ^	Active Queries
616	admin	192.168.1.102	2020-Aug-20 18:15:47	0
646	admin	192.168.1.102	2020-Aug-20 18:15:48	0
908	admin	192.168.1.102	2020-Aug-20 18:22:38	0
917	admin	192.168.1.102	2020-Aug-20 18:22:38	0
1005	escalateduser	192.168.1.102	2020-Aug-20 18:47:18	0

RSA NETWITNESS PLATFORM 11.5.0.0

This is an example of the Services System view for a Log Decoder.

Log Decoder Service Information

Name	LH (Log Decoder)
Version	11.5.0.0 (Rev null)
Memory Usage	2050 MB (6.37% of 32174 MB)
CPU	2%
Running Since	2020-Aug-20 18:14:39
Uptime	3 hours 29 minutes 43 seconds
Current Time	2020-Aug-20 21:44:22

Appliance Service Information

Name	LH (Host)
Version	11.5.0.0 (Rev null)
Memory Usage	31420 KB (0.10% of 32174 MB)
CPU	2%
Running Since	2020-Aug-20 18:14:31
Uptime	3 hours 29 minutes 51 seconds
Current Time	2020-Aug-20 21:44:22

Log Decoder User Information

Name	admin
Groups	Administrators
Roles	aggregate, connections.manage, database.manage, decoder.manage, dpo.manage, index.manage, logs.manage, parsers.manage, rules.manage, sdk.content, sdk.manage, sdk.meta, sdk.packets, services.manage, storedproc.execute, storedproc.manage, sys.manage, users.manage

Host User Information

Name	admin
Groups	Administrators
Roles	appliance.manage, connections.manage, logs.manage, services.manage, storedproc.execute, storedproc.manage, sys.manage, users.manage

Session Information

Session	User	IP Address	Login Time ^	Active Queries
673	admin	127.0.0.1:51258	2020-Aug-20 18:17:12	0
702	admin	127.0.0.1:51362	2020-Aug-20 18:17:12	0
777	admin	127.0.0.1:50378	2020-Aug-20 18:17:30	0
786	admin	127.0.0.1:50380	2020-Aug-20 18:17:31	0
815	admin	127.0.0.1:51436	2020-Aug-20 18:17:31	0

Service Info Toolbar

These two toolbars illustrate the options specific to Decoders and Log Decoders.

In addition to the common options in the Services System view toolbar, you can start and stop capture of packets or logs. The upload file options are different for the standard Decoder (packet capture file) and the Log Decoder (log file).

Action	Description
Upload Packet Capture File	Displays a dialog that provides a way to select a packet capture (.pcap) file for upload to the selected Decoder. For more information, see Upload a Packet Capture File. Note: This option does not apply to Log Decoders.
Upload Log File	Displays a dialog that provides a way to select a log (.log) file for upload to the selected Log Decoder. For more information, see Upload a Log File to a Log Decoder.
Start/Stop Capture	Starts packet capture on the selected Decoder. When packet capture is in progress, the option in the toolbar changes to Stop Capture, and the option to upload a file is unavailable.