NetWitness® Platform

Version 12.4

Decoder Configuration Guide

NETWITNESS

Platform

Contact Information

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

Trademarks

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.

It is advised not to deploy third-party repos or perform any change to the underlying NetWitness Operating System that is not part of the supported NetWitness version. Any such change outside of the NetWitness approved image may result in a service or functionality conflict and require a reimage of the NetWitness system to bring NetWitness back to an optimized functional state. In the event a third-party repo is deployed, or other non-supported change is made by the customer without NetWitness approval, the customer takes full responsibility for any system malfunction until the issue can be remediated through troubleshooting efforts or a reimage of the service.

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 NetWitness Community. 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.

Miscellaneous

This product, this software, the associated documentations as well as the contents are subject to NetWitness' standard Terms and Conditions in effect as of the issuance date of this documentation and which can be found at https://www.netwitness.com/standard-form-agreements/.

© 2024 RSA Security LLC or its affiliates. All Rights Reserved.

April, 2024

Contents

Decoder and Log Decoder Quick Setup **14**

Configure Common Settings on a Decoder **17**

Configure Capture Settings 19

Select a Network Adapter 19

Configure a Decoder to Begin Capturing Data Automatically 21

Configure Optional Capture Settings 22

(Optional) Configure System-Level (BPF) Packet Filtering 24

Examples 26

Testing 26

Conversions 27

(Optional) Configure a Decoder to Capture Data Across All Types of Network Interfaces 27

(Optional) Configure Meta-Only Decoders 30

(Optional) Configure Selective Network Data Collection 31

Use Predefined Policies 32

Create New Policies from Predefined Ones 36

Create Custom Policies 40

Verify Published Policies 44

Troubleshooting Policy Deployments That Fail 52

Unpublish Policies 53

Delete Policies 55

Supported Protocols for Selective Network Data Collection 56

(Optional) Configure a Decoder to Write Standard pcap-formatted Files 58

(Optional) Multiple Adapter Packet Capture 59

Configure Multiple Adapter Packet Capture 59

Per Interface Configuration 60

Stat Nodes 60

Per-Interface Stats 60

Aggregate Stats 61

Pool Stats 61

capture.port Meta Key 61

Special Devices 62

Packet arrived out of order Message 62

Multi-Assembler Modes 62

(Optional) Internet Content Adaptation Protocol Capture 64

ICAP Capture Options 64

(Optional) Data Plane Development Kit Packet Capture 64

Advantages ... 64
How it Works ... 65
Migrate PF_RING Devices to DPDK ... 65
Manually Move a Capture Interface to DPDK ... 66
Check the device status and stats of DPDK interfaces ... 68
Check link status and stats from NwConsole ... 69
Utilizing Receive Side Scaling with DPDK ... 70
Advantages ... 70
Configuring Receive Side Scaling ... 71
How it Works ... 71
Configure the Decoder Capture Failover in Load Balance Deployments ... 71
Configuration Failover on DPDK Capture Interfaces ... 71
Disable Wake On LAN (WOL) in BIOS settings ... 72
(Optional) Remove DPDK on Capture Interfaces ... 72
Validate Link Status ... 72
DPDK Interface Configuration and Statistics ... 73
(Optional) Preserve VLAN Tags When Using the Packet MMAP Capture Interface ... 73
(Optional) Process Raw Syslog Data without Priority Field ... 78
(Optional) Configure Decoder to Support OpenAppID ... 79
Enable and Disable Parsers and Log Parsers ... 82
(Optional) Decoder Parser Utilization Insights ... 84
Configuration and Usage ... 84
Start and Stop Data Capture ... 85
Configure Decoder Rules ... 86
Rule Processing ... 86
Rule and Query Guidelines ... 87
Rule Examples ... 87
Invalid Rules ... 88
General Syntax Guidelines ... 88
Capture Rule Syntax ... 89
Configure Capture Rules ... 91
Import Rules from a File and Export Rules ... 93
Push Rules to Other Services ... 95
Change Execution Order of Rules ... 97
Restore a Rule Snapshot from History ... 97
Configure Application Rules ... 99
Configure Correlation Rules ... 103
Configure Network Rules ... 107
Supported Meta Keys in Network Rule Conditions ... 107
Fix Rules with Invalid Syntax ... 111
Decoder Commands for Managing Rules ... 113

add Command ... 113
merge Command ... 114
Methods of Sending a List of Rules to a Service ... 114
Send a NetWitness Rule File ... 115
Send Numbered Parameters ... 115
Ordering Rules When Pushing ... 116
replace Command ... 117
clear Command ... 117
delete Command ... 117
validate Command ... 117
Configure Parsers and Feeds ... 118
Configure Parsers ... 118
Upload and Delete Custom Parsers ... 119
Upload Parsers to a Decoder or Log Decoder ... 119
Manage Upload Jobs ... 121
Delete Deployed Parsers ... 121
Enable and Configure the Entropy Parser ... 122
Entropy Parser Configuration in the Concentrator Custom Index File ... 125
Flex Parsers ... 127
Arithmetic Functions ... 128
Language Definition ... 128
Common Parser Operations ... 130
Match Port and Identify Immediately ... 130
Match Port and Delay Identification ... 130
Match Token and Identify Immediately ... 131
Match Multiple Tokens ... 131
Match Token and Create Metadata ... 132
General Functions ... 133
General Functions Language Definition ... 133
Logging Functions ... 135
Language Definition ... 135
Nodes ... 136
Nodes Language Definition ... 136
Payload Functions ... 140
Language Definition ... 140
Regex ... 142
Language Definition ... 142
String Functions ... 143
String Functions Language Definition ... 143
GeoIP2 Parsers ... 146
Lua Parsers ... 148

HTTP Parsers 149
Visibility into HTTP/2 Sessions 150
To turn on header parsing and decompression for HTTP/2 sessions: 150
To avoid duplicate meta from HTTP/2 parser when Lua parser HTTP_lua is enabled: 150
Decoder Snort Detection 152
Using Snort Rules in Decoders 152
Initial Decoder Snort Configuration 152
Snort Parser Output 154
Rule Statistics 155
Configuration Directives 155
Variable Definitions 155
ruletype 156
Other Configuration Entries 156
config detection: debug 156
config nopcre 156
config classification 156
Snort Parser Capabilities 156
Rule Sections 156
Content 157
Flow 158
Fast Pattern 158
PCRE 158
HTTP Parsing 158
File Data 159
Binary (byte) Directives 159
Default option / fallback 160
Packet Scanning/Scoping 160
Known Limitations for Snort Parser 160
Decoder Token Match Limits 160
Flow Bits 160
DCE Preprocessor 160
Thresholds and Filters 160
Raw Byte Scanning 160
IPv6 161
Performance Considerations 161
Content Tokens 161
Negation 161
Search Parser 162
Search Methods 162
Syntax 162
162
Wireless LAN Configuration 163
Troubleshooting Parsers 163

Lua Parser Errors 163
Results of Automatically Disabling a Parser 164
Resetting the Parser 164
Configure Feeds 164
Custom Feed Definition File Structure 165
Sample Feed Definition File 165
Define Multiple Values in a Single Field 166
Feed Definition Equivalents for Custom Feed Wizard Parameters 166
Sample Files for a MetaCallback Feed Using CIDR Index Range for IPv4 and IPv6 168
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. 168
The following example shows the contents of both a .csv file and an .xml file for a MetaCallback feed using CIDR index ranges for IPv6. 169
Feed Definitions File 170
Create a Custom Feed 171
Create a STIX Custom Feed 181
Create an Identity Feed 187
Import the SSL Certificate 194
Cannot Verify Identity Feed URL 194
Edit, Upload, or Remove a Feed 196
Create Custom Meta Keys Using a Custom Feed 200
Add a Custom Meta Key in the Log Decoder 200
Deploy a Log Decoder Feed in Live 201
Add the Custom Meta Key Entry in the Concentrator Custom Index file 205
Investigate a Custom Meta Key 205
Additional Procedures 206
Verify the Custom Meta Keys on ESA 206
Update the Schema in Archiver 207
Update the Schema in Warehouse Connector 207
Update the Schema in Reporting Engine 207
Decoder and Log Decoder Additional Procedures 209
Configure High Speed Packet Capture Capability 209
Hardware Prerequisites 211
Software Prerequisites 211
Decoder Installation 211
Install the Decoder Service on Host that Will Perform the Capture 211
Select the Capture Interfaces and Assign Interfaces to DPDK 211
Verify the Decoder Packages Installation 212
Select the Decoder Operating Mode 212
Configure the Decoder for Normal Mode (Default Mode) 212
Configure the Decoder for 10G Mode 212

Performance Tuning Parameters 214
Storage Considerations 215
Packet Retention Requires Extremely High Sustained Throughput 215
Using SAN and Other Storage Configurations 215
Optimize Read/Write Operations When Adding New Storage 215
Parsing and Content Considerations 217
Best Practices 217
Use Case 1: 10G Mode - Egress, Deep Packet Inspection 218
Tested Live Content 218
Not Tested 218
Other 218
Aggregation Adjustments Based on Tested Live Content 219
Configure the Decoder for NDR Mode 219
Performance Tuning Parameters 221
Storage Considerations 221
NDR Mode Assumes No Packet Retention 221
Parsing and Content Considerations 222
Parsing at Speeds Greater than 10G 222
Best Practices 222
Use Case 1: NDR Mode - Egress, General Purpose 223
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + All Native Parsers 223
Tested Live Content 223
Not Tested 223
Other 223
Use Case 2: NDR Mode - Egress, Data Exfiltration 224
Generate NetFlow Style Meta Only + Exfiltration Focused Specific Native Parsers 224
Tested Live Content 224
Not Tested 224
Other 224
Use Case 3: NDR Mode - Lateral Movement 224
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + Specific Native Parsers 224
Tested Live Content 224
Not Tested 225
Other 225
Configure High Speed Packet Capture Capability 226
Hardware Prerequisites 227
Software Prerequisites 227
Decoder Installation 227
Install the Decoder Service on Host that Will Perform the Capture 228
Select the Capture Interfaces and Assign Interfaces to DPDK 228
Verify the Decoder Packages Installation 228
Select the Decoder Operating Mode 228

Configure the Decoder for Normal Mode (Default Mode) 229
Configure the Decoder for 10G Mode 229
Performance Tuning Parameters 230
Storage Considerations 231
Packet Retention Requires Extremely High Sustained Throughput 231
Using SAN and Other Storage Configurations 231
Optimize Read/Write Operations When Adding New Storage 231
Parsing and Content Considerations 233
Best Practices 233
Use Case 1: 10G Mode - Egress, Deep Packet Inspection 234
Tested Live Content 234
Not Tested 235
Other 235
Aggregation Adjustments Based on Tested Live Content 235
Configure the Decoder for NDR Mode 235
Performance Tuning Parameters 237
Storage Considerations 238
NDR Mode Assumes No Packet Retention 238
Parsing and Content Considerations 238
Parsing at Speeds Greater than 10G 238
Best Practices 238
Use Case 1: NDR Mode - Egress, General Purpose 239
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + All Native Parsers 239
Tested Live Content 239
Not Tested 239
Other 240
Use Case 2: NDR Mode - Egress, Data Exfiltration 240
Generate NetFlow Style Meta Only + Exfiltration Focused Specific Native Parsers 240
Tested Live Content 240
Not Tested 240
Other 240
Use Case 3: NDR Mode - Lateral Movement 240
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + Specific Native Parsers 240
Tested Live Content 241
Not Tested 241
Other 241
Configure a Log Decoder to Accept Protobuf 242
Configure Session Split Timeouts 244
Configure Syslog Forwarding to Destination 247
Configure Transaction Handling on a Decoder 249
Transactions Off 249
Transactions Represented as Meta Items 250

Transactions Split Sessions 250
Configure Data Export 251
Decrypt Incoming Packets (TLS 1.2) 252
Step 1: Validate That The Network Decoder Captures Encrypted Traffic 254
Step 2: Obtain Private Keys from Managed Servers 254
Extracting Certificates for the Decoder SSL Decryption 254
Extract private key from Apache 254
Extract Private Key and Certificate from Windows Servers 255
Step 2 (a): Extract PKCS12 file from Internet Information Services (IIS) 255
Using IIS Manager App 255
Using Microsoft Management Console (MMC) 255
Step 2 (b): Extract PKCS12 file from Exchange Servers 255
Step 2 (c): Extract Private Key From a pfx File Using OpenSSL 256
Recommendations 256
Step 3: Validate That The Private Key Cipher Suite is Supported 256
Supported Cipher Suites 256
Unsupported Cipher Suites 257
Step 4: Confirm HTTPS Parser is Enabled on Decoders 261
Step 5: Upload the Supported Private Keys to Decoders 261
Upload Multiple Premaster and Private Keys 263
Parameters for Managing Keys 266
Return Values 267
Encryption Keys 267
Premaster Key 267
Private Keys or PEM Files 267
Step 6: Reload HTTPS Parser on Decoders 268
Step 7: View Decrypted Sessions 269
Viewing Unencrypted Traffic 269
Troubleshooting 270
Decryption of Secure SMTP 270
Performance Considerations 271
TLS Certificate Hashing 273
JA3 and JA3S TLS Fingerprints 274
Perform Simultaneous Ingestion of the Encrypted and Decrypted Traffic Streams to Decoder 274
Community ID 275
Decrypt Incoming Packets TLS 1.3 276
TLS 1.3 Keys and Cipher suites 276
Cipher Suites 276
Steps to configure Network Decoder to decrypt TLS 1.3 traffic 277
Step 1: Validate that the Network Decoder Captures Encrypted TLS 1.3 Traffic 277

Step 2: Confirm HTTPS Parser is Enabled on Decoders 277
Step 3: Integrate Network Decoder to receive keys from Thirdparty TLS 1.3 key forwarder 278
Option 1: Configure F5 BIG-IP Local Traffic Manager (LTM) 278
Option 2: Using sslKeys API 278
Parameters for Managing Keys 279
Usage: 280
Return Values 280
Step 4: (Optional) Upload TLS 1.3 keys to validate Decryption 280
TLS 1.3 Keys 281
Step 5: Validate TLS 1.3 keys are received on Decoder 281
Step 6: View Decrypted Sessions 282
Viewing Unencrypted Traffic 282
Option 1 : Using Event Analysis UI 282
Option 2 : Using RESTful API on Decoder service 282
Troubleshooting 283
Additional Features 283
Edit Decoder System Configuration 284
Enable CPU Usage Statistics for Installed Content 286
Enable Parser Mappings 287
Enable IP Address to Event Source Mapping 287
Update IP to Event Source Mapping 288
Read IP to Event Source Type Mappings 290
Edit IP to Event Source Type Mappings 290
Delete IP to Event Source Type Mappings 291
Sort the Hostname or Event Source Type 291
Import IP to Event Source Mapping Entries 292
Export IP to Event Source Mapping Entries 292
Search IP to Event Source Mapping Entries 293
Enable or Disable Lua and Flex Parsing Systems 294
Map IP Address to Service Type for Log Parsing 296
Map an IP Address to a Service Type 296
IPdevice Command 296
Map an IP Address to a Time Zone 298
tzinfo Command 298
iptmzone Command 299
Examples 299
Change the Date format 299
Missing Year Support 301
Latent log during transition to new year 302
Logs from forward time zones (or skewed clocks) just before new year transition 302

Latent logs received where a leap day cannot be successfully assigned to an appropriate leap year ... 302
Limitations ... 302
Examples ... 302
Obtain Log Files from a Log Decoder ... 304
Upload a Log File to a Log Decoder ... 306
Upload a Packet Capture File ... 307
Simultaneous Import and Capture ... 308
F5 BIG IP - NetWitness Perfect Forward Secrecy Inspection Visibility ... 309
Prerequisites ... 309
Deployment ... 310
BIG-IP LTM Configurations ... 310
HTTPS Monitor ... 311
Pool & Node ... 313
HTTP Profile ... 314
Server SSL Profile ... 316
Internal Virtual Server ... 318
iRule Session keys copying and Sideband communication ... 320
Configuration ... 320
iRule NetWitness-sideband-irule ... 320
Troubleshooting Packet Drops ... 325
Quick configuration Checks to avoid packet drops ... 325
To check and tune the configuration: ... 325
Information required to troubleshoot packet drops ... 326
What if the REST port is inaccessible? ... 326
How do I troubleshoot packet drops? ... 327
Navigate to the Packet drops tool using the REST interface ... 327
Introduction to Drops tool Charts ... 327
Symptoms Checklist ... 328
Symptom 1: Higher traffic ingestion rates for the content deployed would cause packet drops ... 328
Resolution: ... 328
Symptom 2: Packet Database Write backup would cause packet drops ... 329
Resolution: ... 330
Symptom: Session Database Write backup would cause packet drops ... 333
Resolution: ... 334
Symptom: Session Pool Parsing delays causes packet drops (Parsing issues) ... 335
Resolution: ... 337
Symptom: Session Parsing stuck can cause packet drops ... 339
Resolution: ... 339
Symptom: Lua Parser Warnings can cause packet drop issues ... 339

Decoder and Log Decoder References **341**
Services Config View - Capture Policies Tab 342
Services Config View - Edit Capture Policies Wizard 345
Services Config View - Data Privacy Tab 349
Services Config View - Data Retention Scheduler 351
Services Config View - Feeds Tab 353
Upload Feeds Dialog 356
Services Config View - Files Tab 358
Services Config View - General Tab 360
Services Config View - Parsers Tab 369
Services Config View - Parser Mappings Tab 371
Data Export Tab 374
Services Config View - Rules Tabs 376
App Rules Tab 380
Correlation Rules Tab 384
Network Rules Tab 387
Services System View - Decoders 390

Decoder and Log Decoder Quick Setup

A basic NetWitness 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 Administration, Live Services, and Investigate. To add a service in NetWitness, 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 services.

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

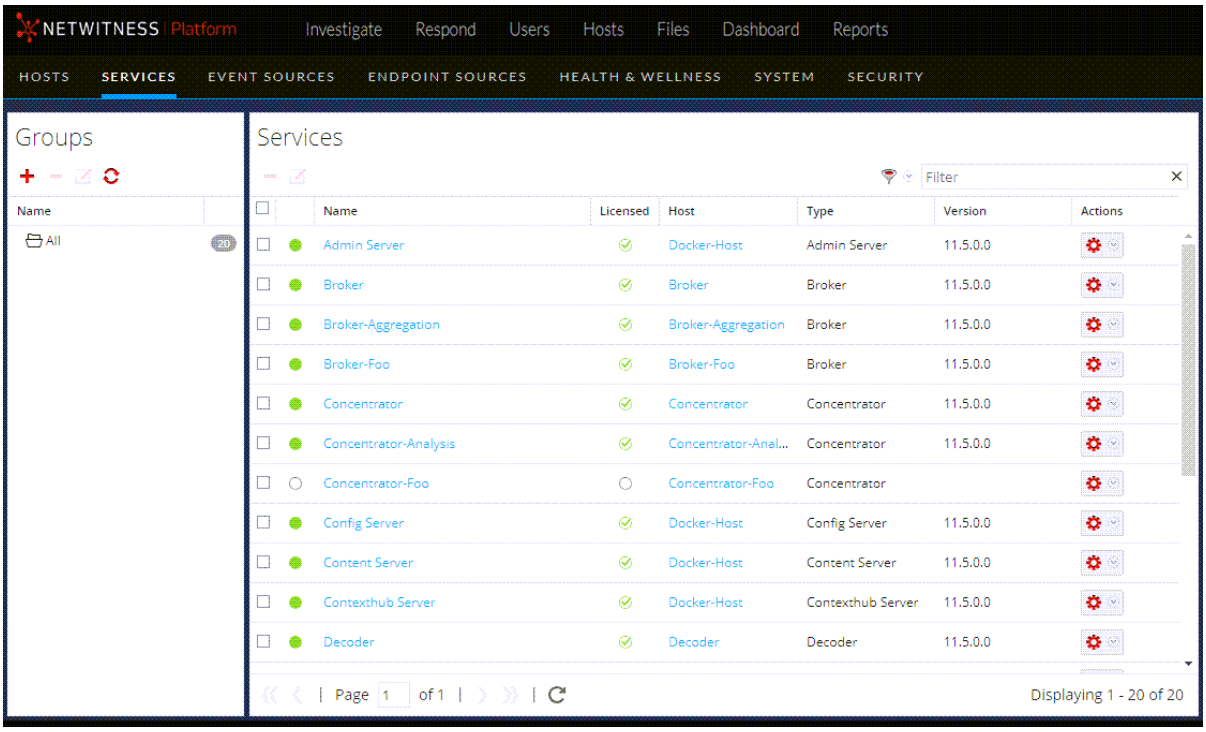
1. Decoders
2. Log Decoders
3. Concentrators (refer to the *Broker and Concentrator Configuration Guide*)
4. Brokers (refer to the *Broker and Concentrator Configuration Guide*)

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.

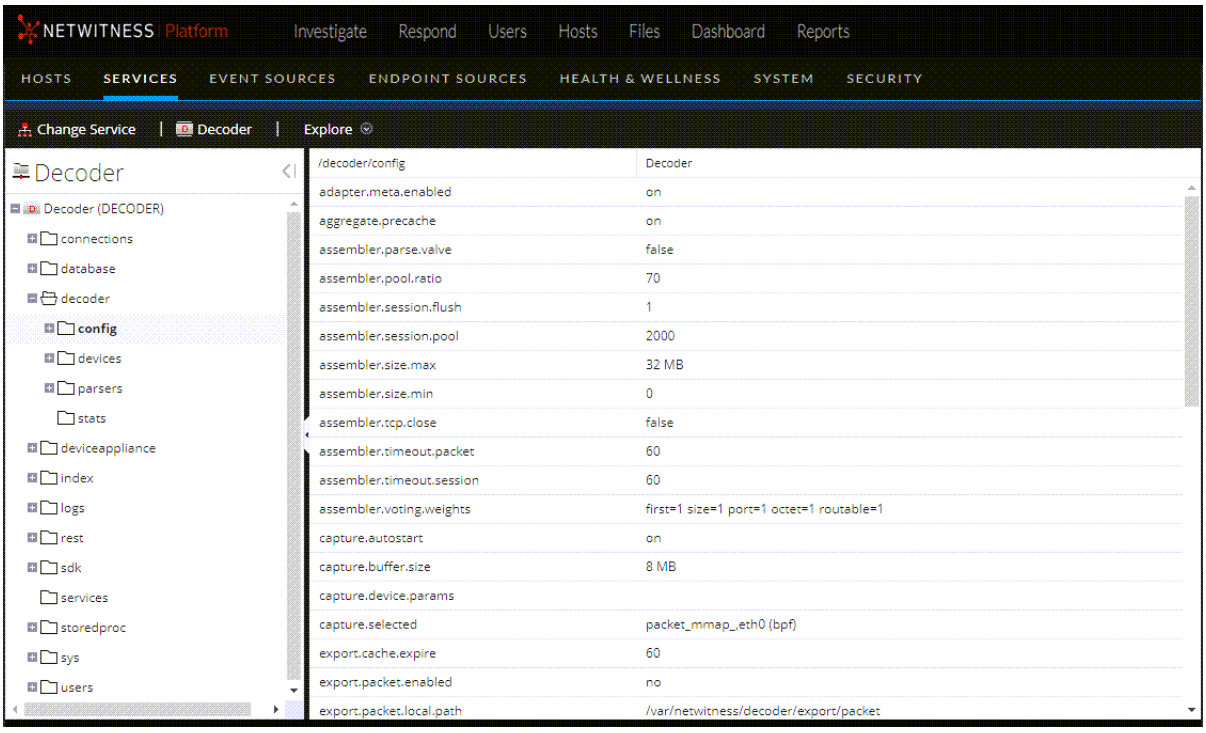
Basic configuration of the Decoder involves selecting a network adapter interface and starting data capture.

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 user interface. For the most part, configuration is performed in the Administration Services view ( (Admin) > Services).



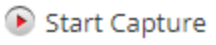
Administrators who feel comfortable working outside of the user interface can configure the basic parameters as well as advanced settings by editing database nodes in the Decoder node tree using the Services Explore view.

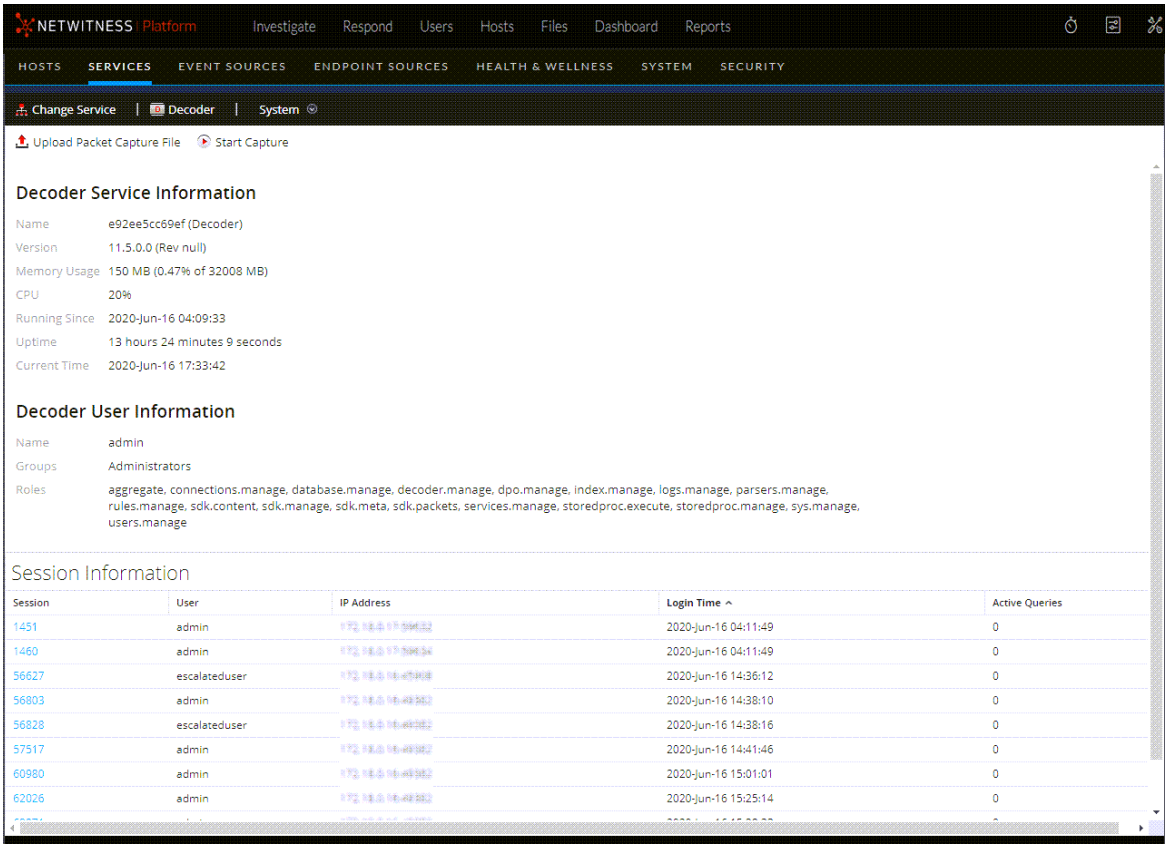


Perform Initial Quick Setup

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  .



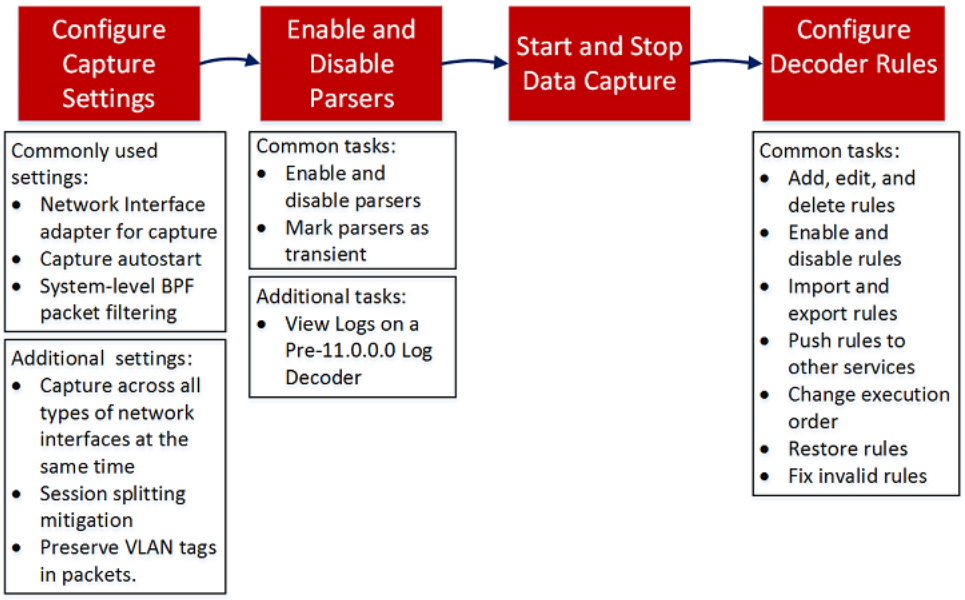
   b. To enable Capture Autostart, see "Configure a Decoder to Begin Capturing Data Automatically" in Configure Capture Settings.
   The Decoder begins capturing data for consumption by a Concentrator. For additional configuration options, refer to Configure Common Settings on a Decoder Configure Common 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.

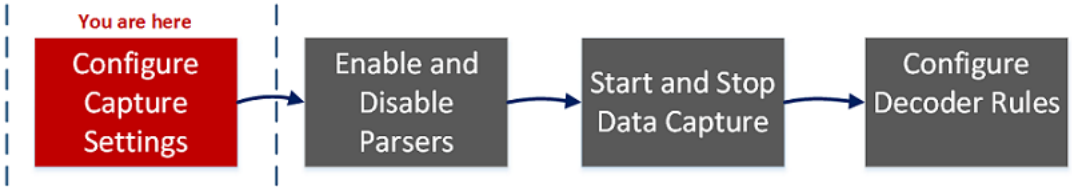


Configuration Step	Description
Configure Capture Settings	When initially setting up the Decoder, configuring the network adapter interface is required. Additional optional capture settings are available; one that is frequently used is Capture Autostart.
Enable and Disable Parsers and Log Parsers	View 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. 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 *Live Services Management Guide*. 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

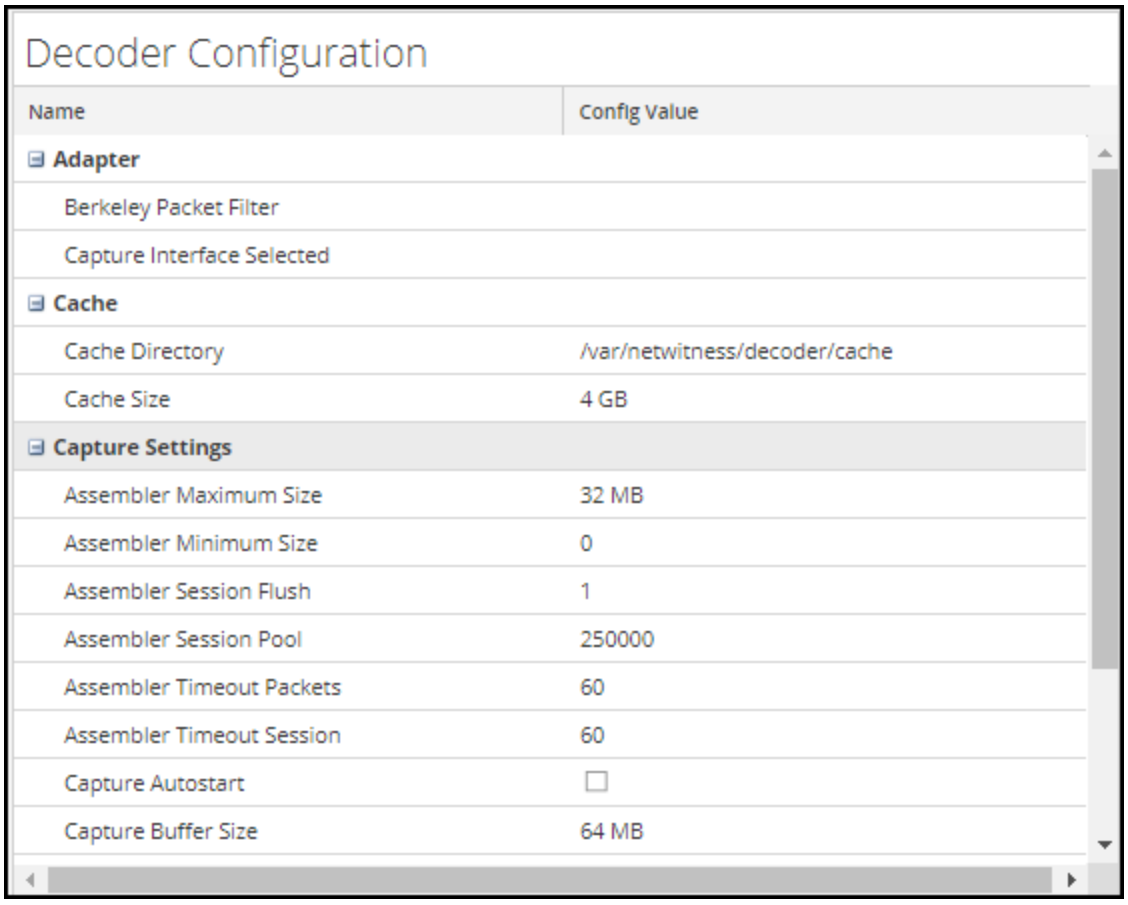
The table below describes the Network Adapter settings for a Decoder. The system administrator sets the default network adapters when the Decoder is installed. Consult your System Administrator for more information.

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 `packet_mmap_,7,eth1` adapter, which corresponds to the monitor port located on the motherboard. There are six additional capture ports: • `packet_mmap_,1,lo (bpf)` • `packet_mmap_,2,eth2 (bpf)` • `packet_mmap_,3,eth3 (bpf)` • `packet_mmap_,4,eth4 (bpf)` • `packet_mmap_,5,eth5 (bpf)` • `packet_mmap_,8,ALL (bpf)` There are three wireless capture services available: • `packet_netmon_ (Microsoft Netmon)` • `packet_mac80211_ (Linux mac80211)` • `packet_airport_ (Mac OS X AirPort)`
Capture Interface Selected for Log Decoder	The following capture service is available: • `log_events,Log Events`

To configure the network adapter on a Decoder:

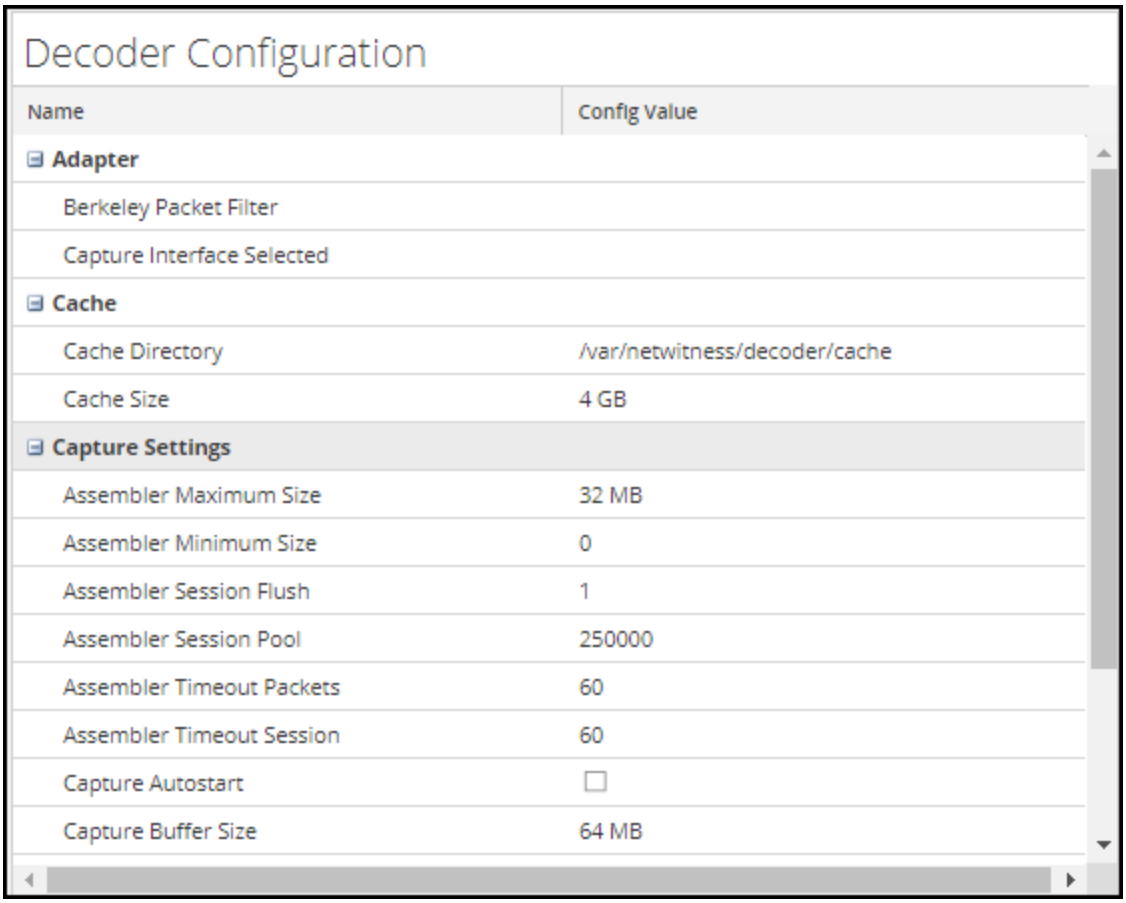
1. Go to  (Admin) **> Services**.
2. In the **Administration Services view**, select the Decoder and  **> View > Config.**
   The Services Config view is displayed with the General tab open.



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

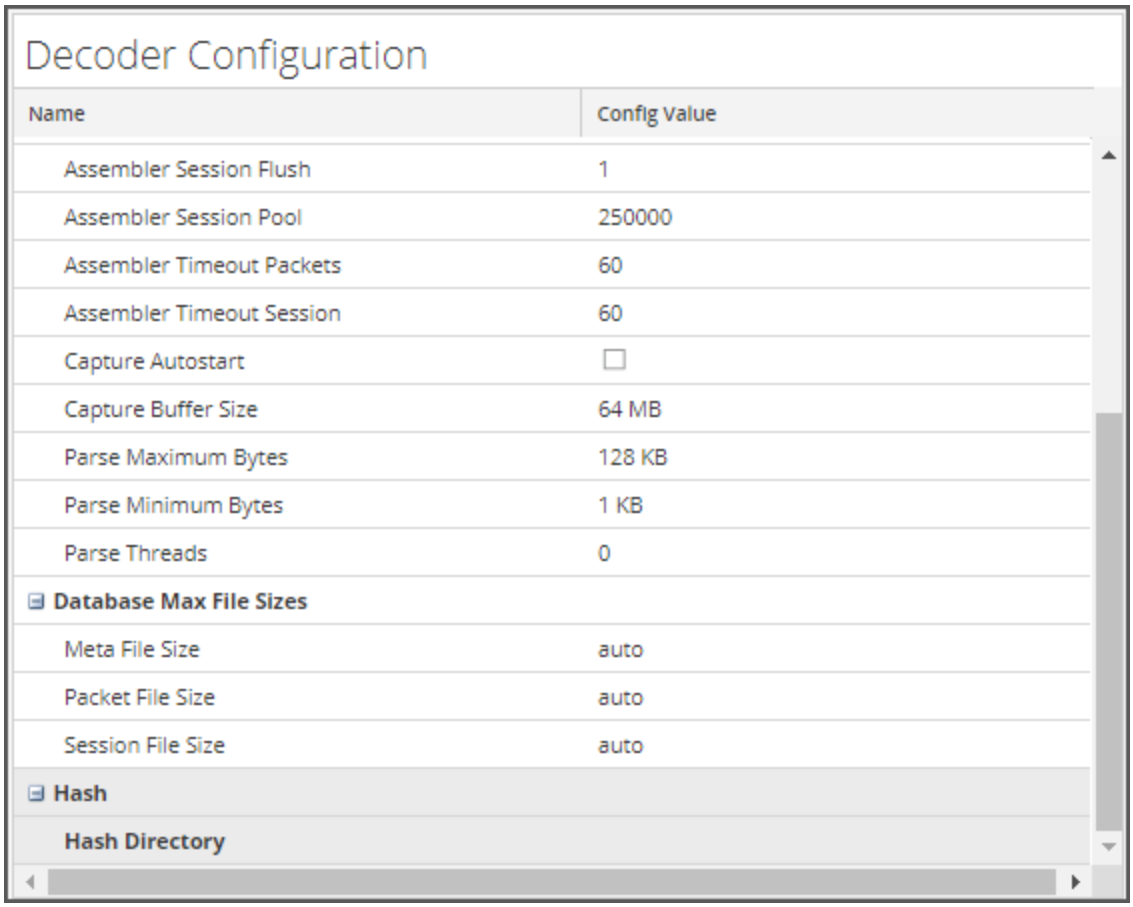
1. Go to (Admin) **> Services**.
2. In the **Administration Services view**, select the Decoder and **> View > Config.**
   The Services Config view is displayed with the General tab open.



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.



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)
```

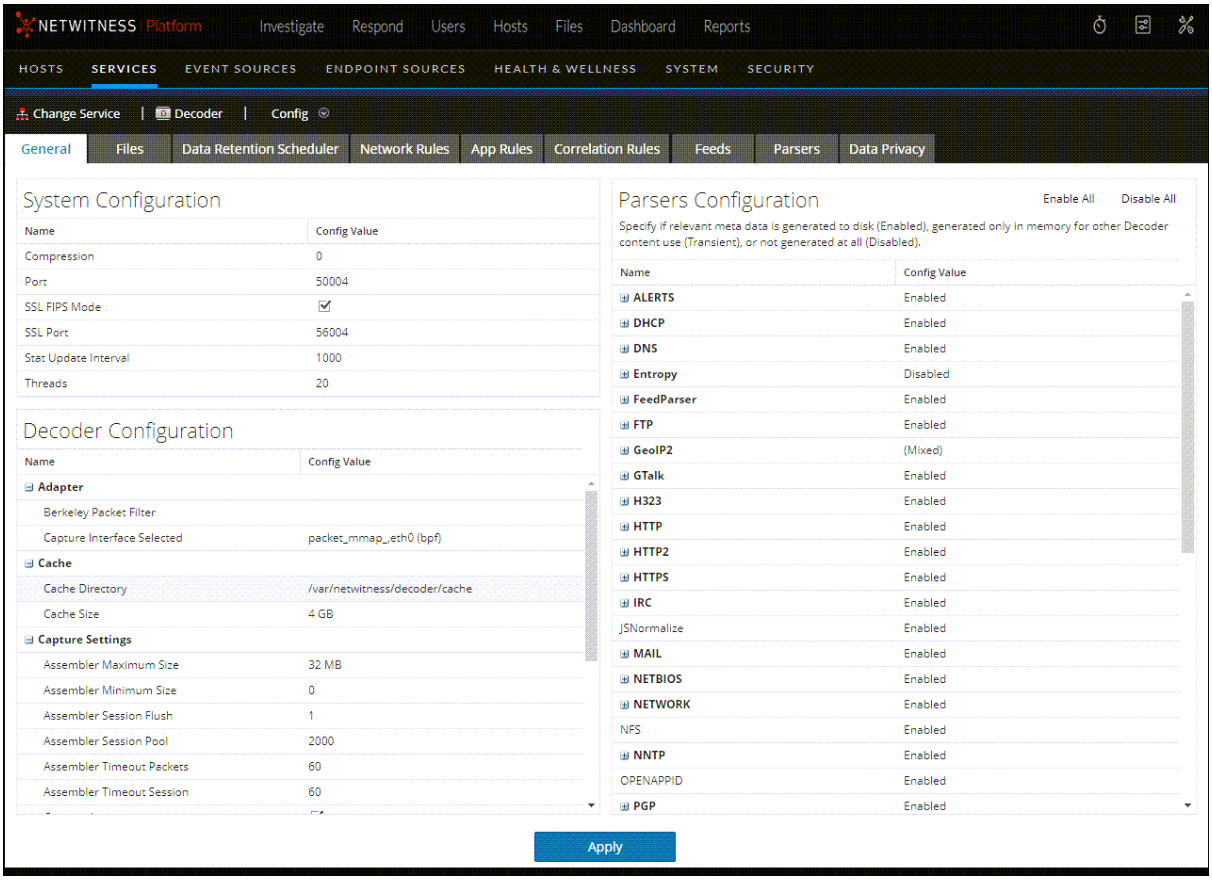
For a full reference of the `Libpcap` filter syntax, see the main pages for:

- 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.**

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



3. In the **Decoder Configuration Section**, under **Adapter**, click in the field next to **Berkeley Packet Filter**.
4. 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.
5. 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
   ```
6. To activate the filter, you must stop and start capture on the Decoder:
   a. Change the **Config** view to the **System** view.
   b. Click **Stop Capture**.
   c. 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 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 `em1`.

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-separate list of interfaces that are used for capture. Instead of using all interfaces for capture, only the specified interfaces are used.

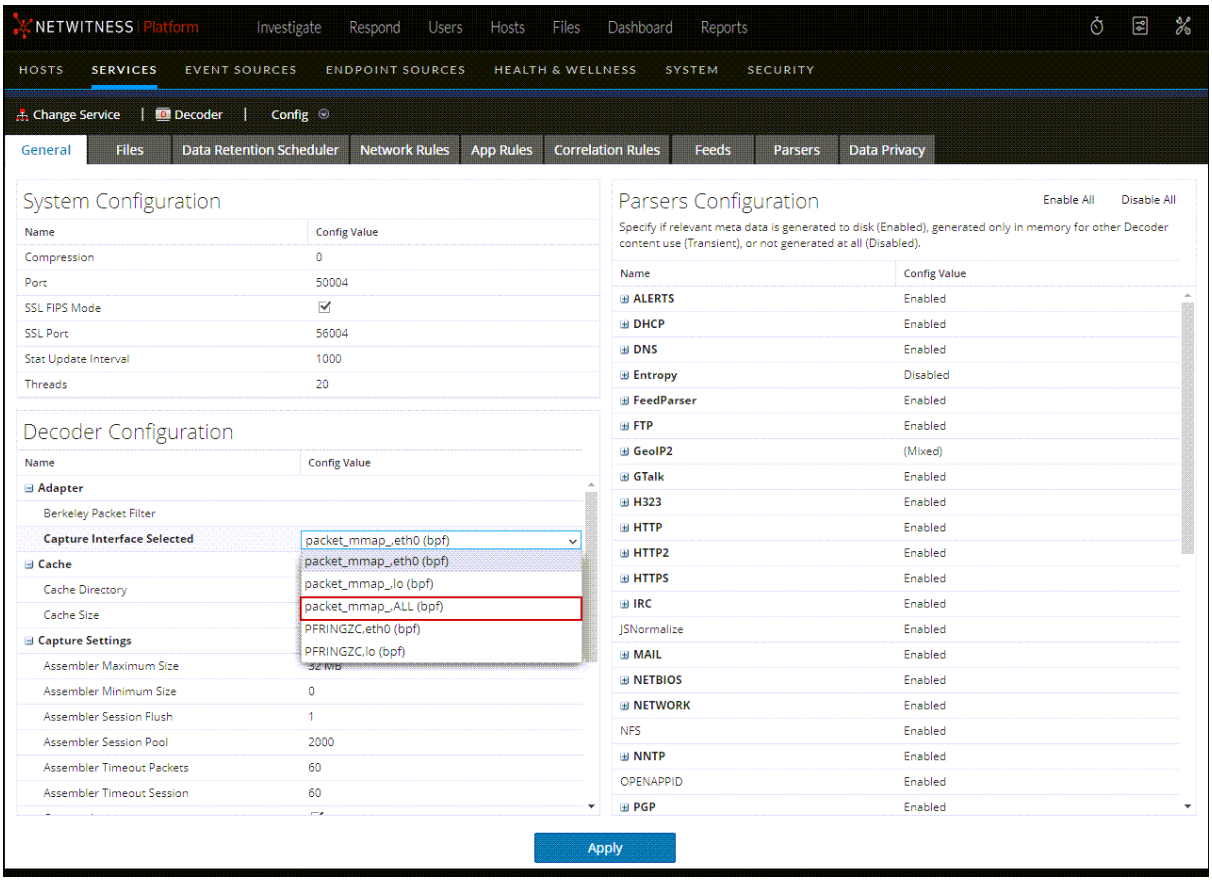
For example, if you want to force capture on interfaces `em1`, `em2`, and `em4`, and ignore `em3`, you can select the `packet_mmap_,ALL` adapter, and then add this line to `capture.device.params`: `interfaces=em1,em2,em4`

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.

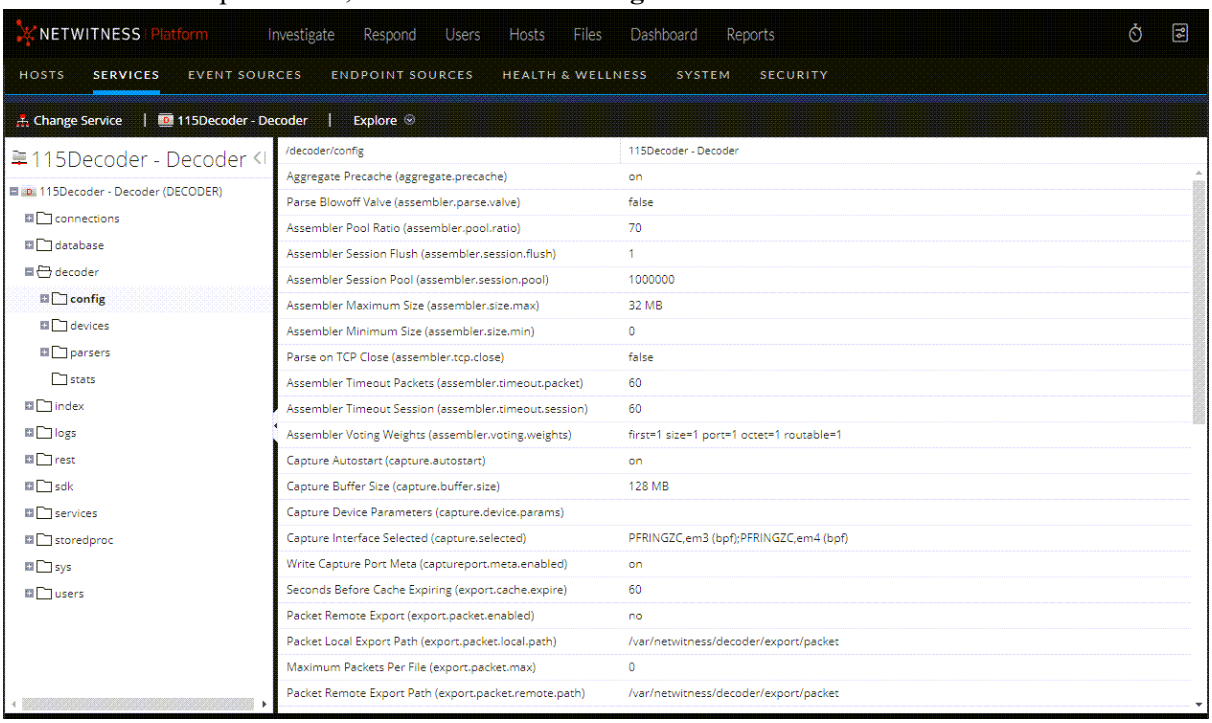
To configure the `packet_mmap_,ALL` adapter to capture from specific interfaces instead of all interfaces:

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.

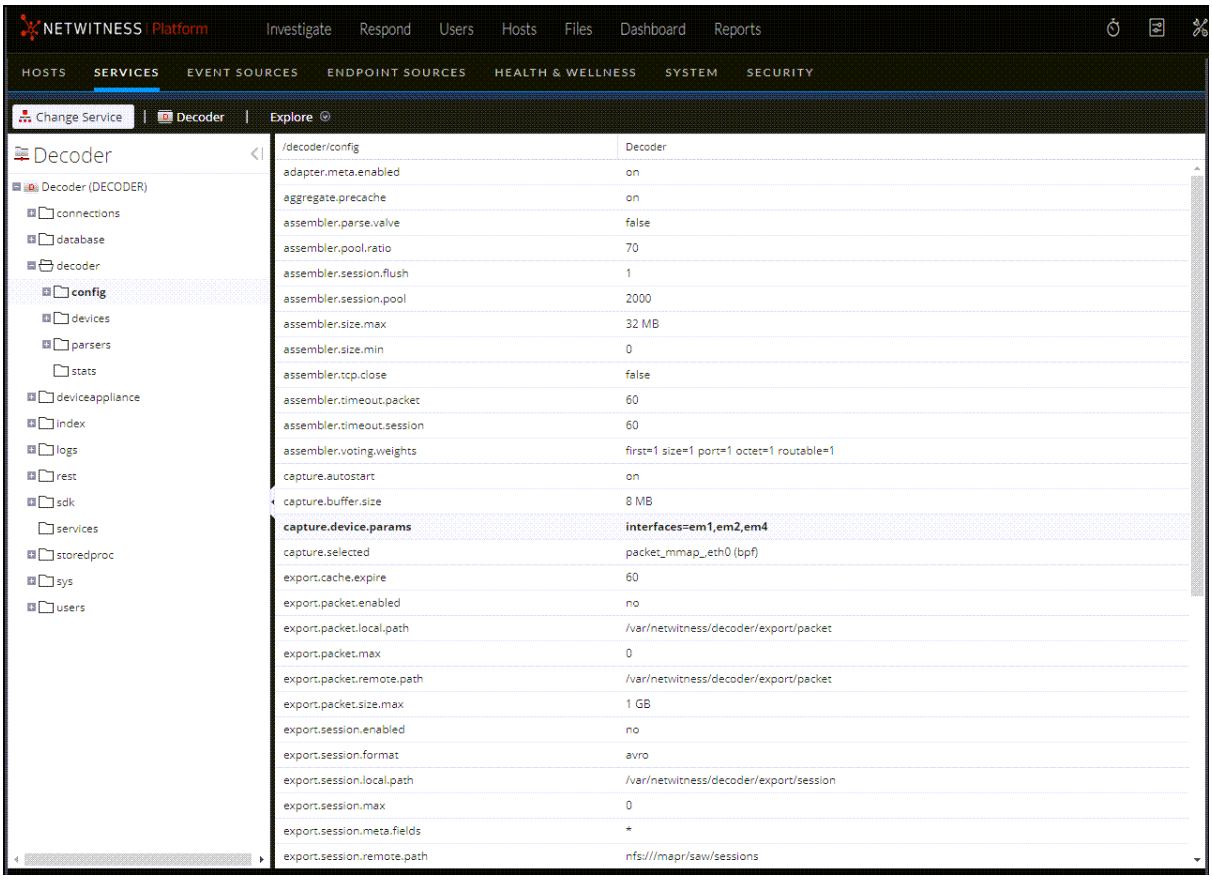


3. Click **Apply**, and then restart the Decoder service.

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



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



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 NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

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 NetWitness 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 polices 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

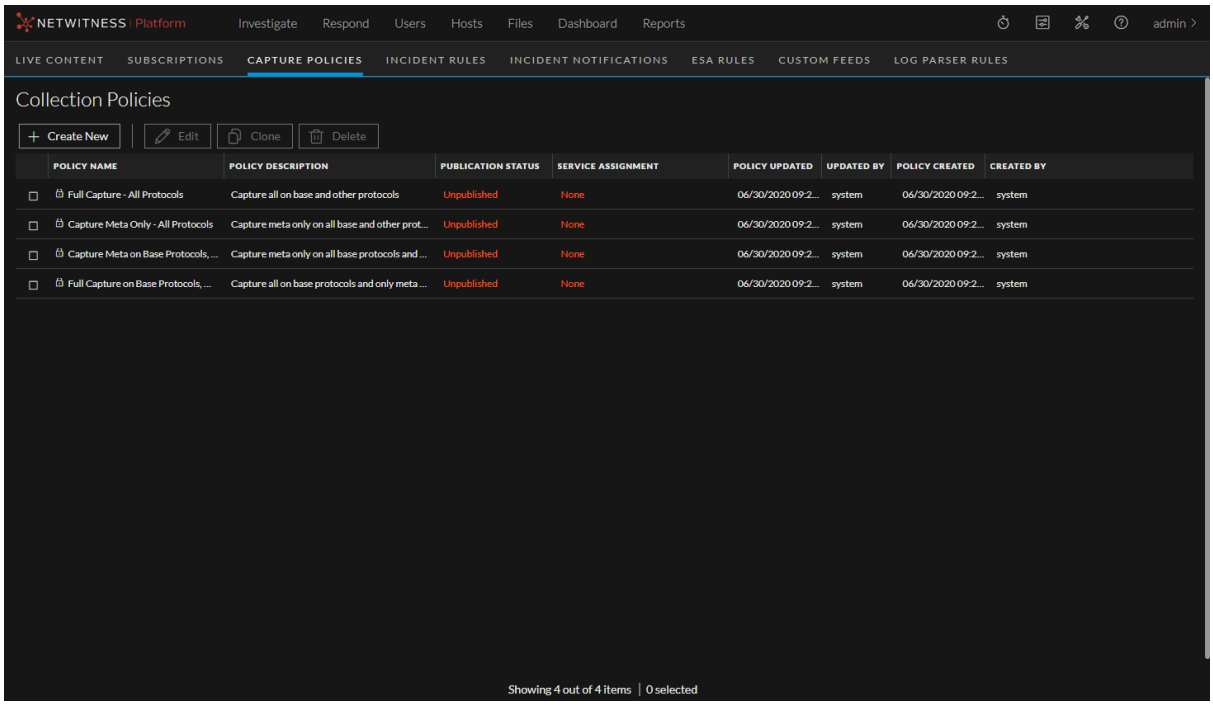
Use Predefined Policies

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.

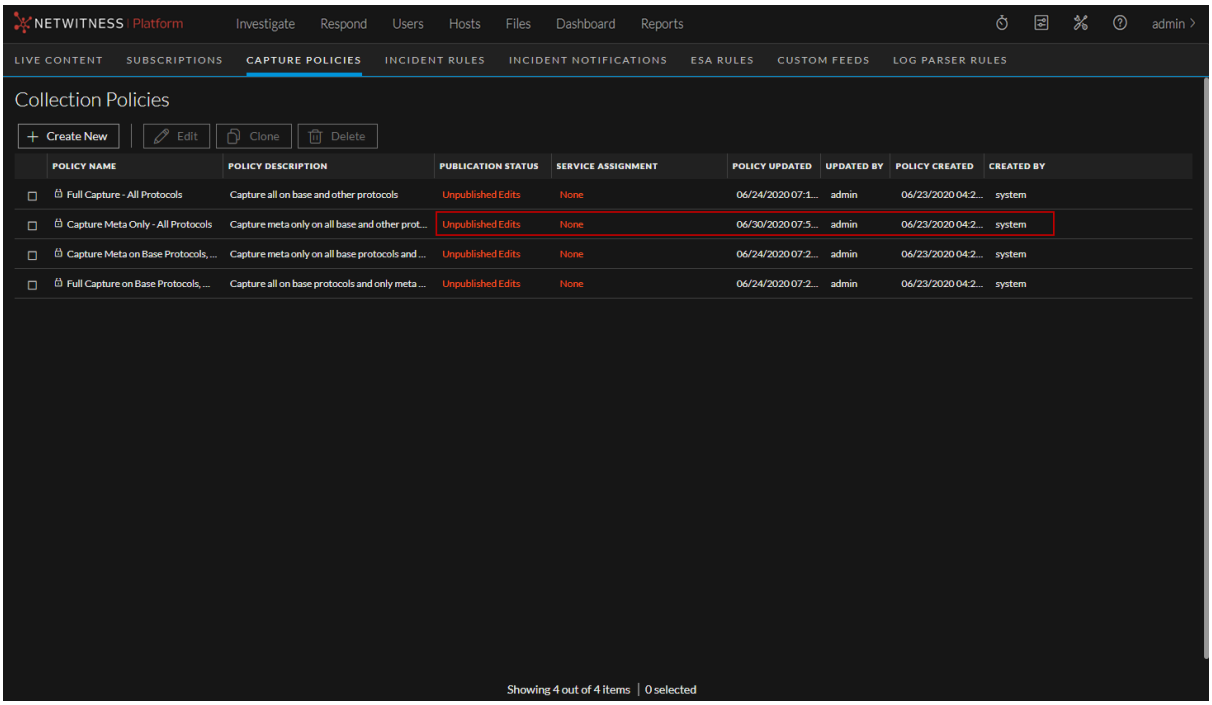
In this situation, the only steps you must complete is to select the "Capture Meta Only – All Protocols" predefined policy to edit it, enable it by assigning it to the appropriate Decoders, and then publish the policy to start collecting the data.

To work with a predefined policy:

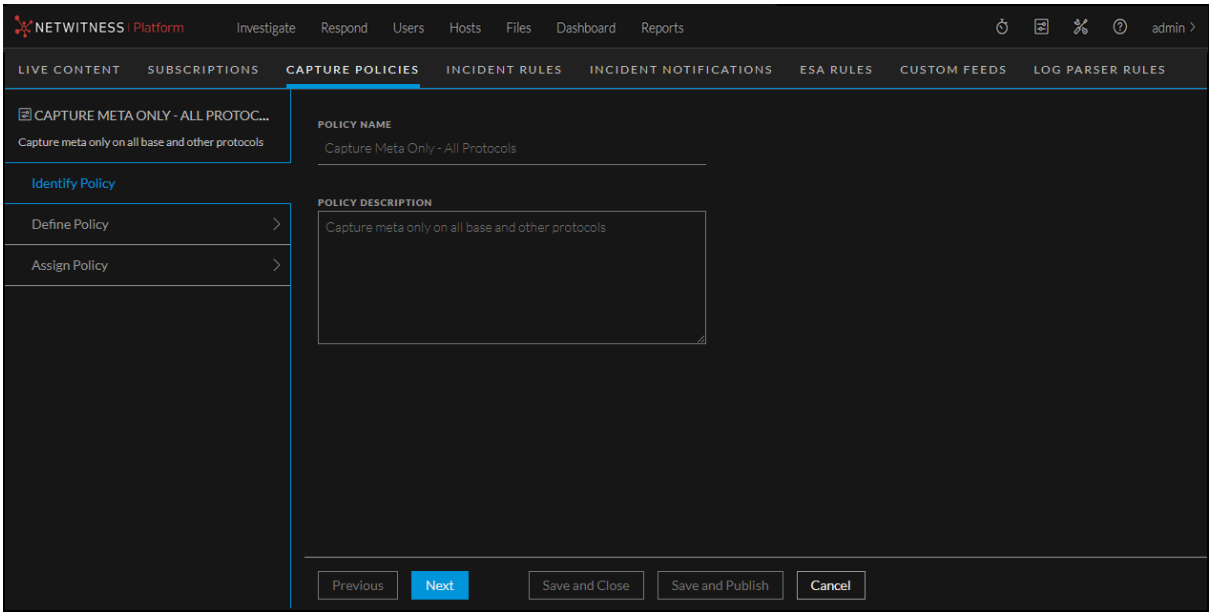
1. In the NetWitness user interface, go to (Configure) and select the **CAPTURE POLICIES** tab.



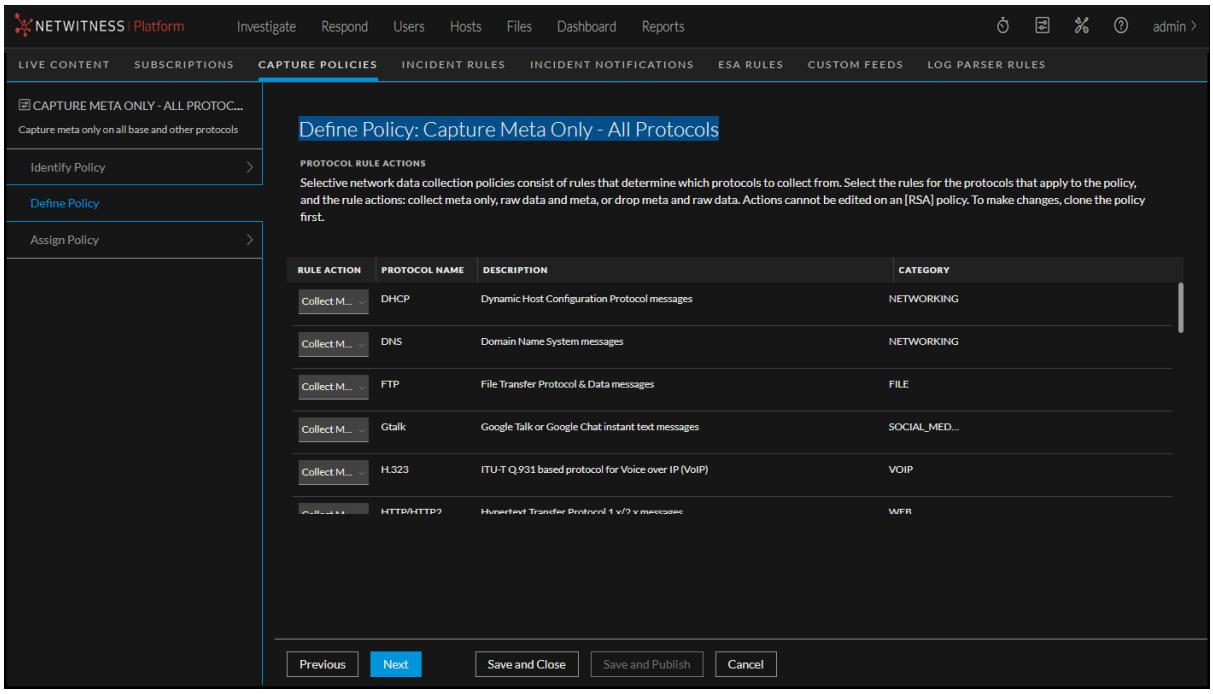
Notice that the PUBLICATION STATUS column lists each of the predefined policies with the status of **Unpublished**. The only way to edit a predefined policy is to assign Decoders to apply the policy. If changes are made to the Decoder assignments but are not published, the status changes to **Unpublished 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.



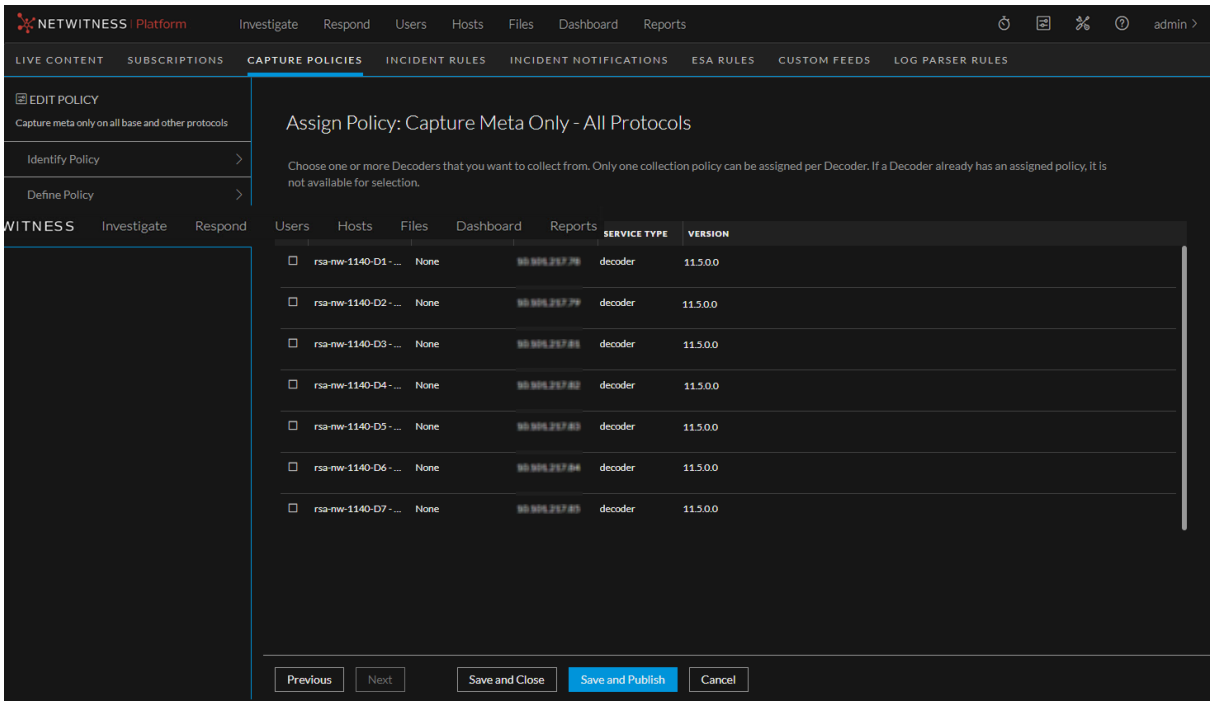
2. Select **Capture Meta Only - All Protocols**, and then click **Edit**.
The editing view is displayed with the **Identify Policy** option active.



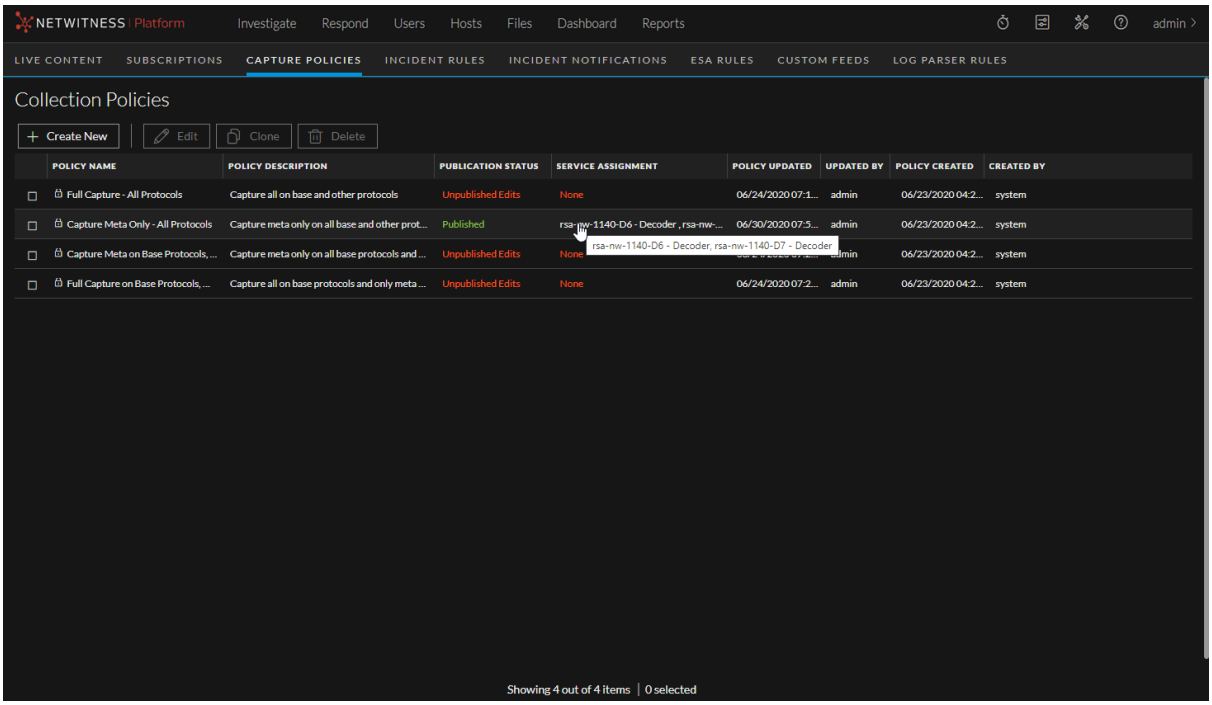
3. Because this is a predefined policy, you cannot change the policy name or description. Click **Next**.
The Define Policy: Capture Meta Only - All Protocols page is displayed.



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.



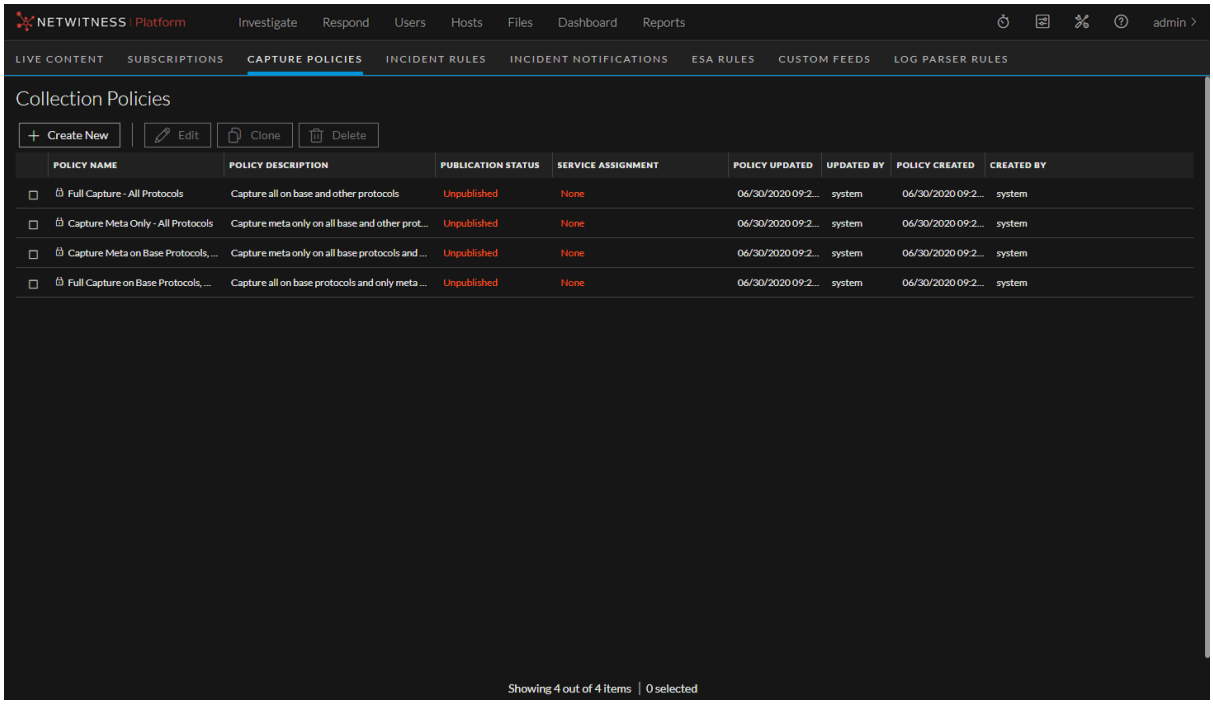
Create New Policies from Predefined Ones

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.

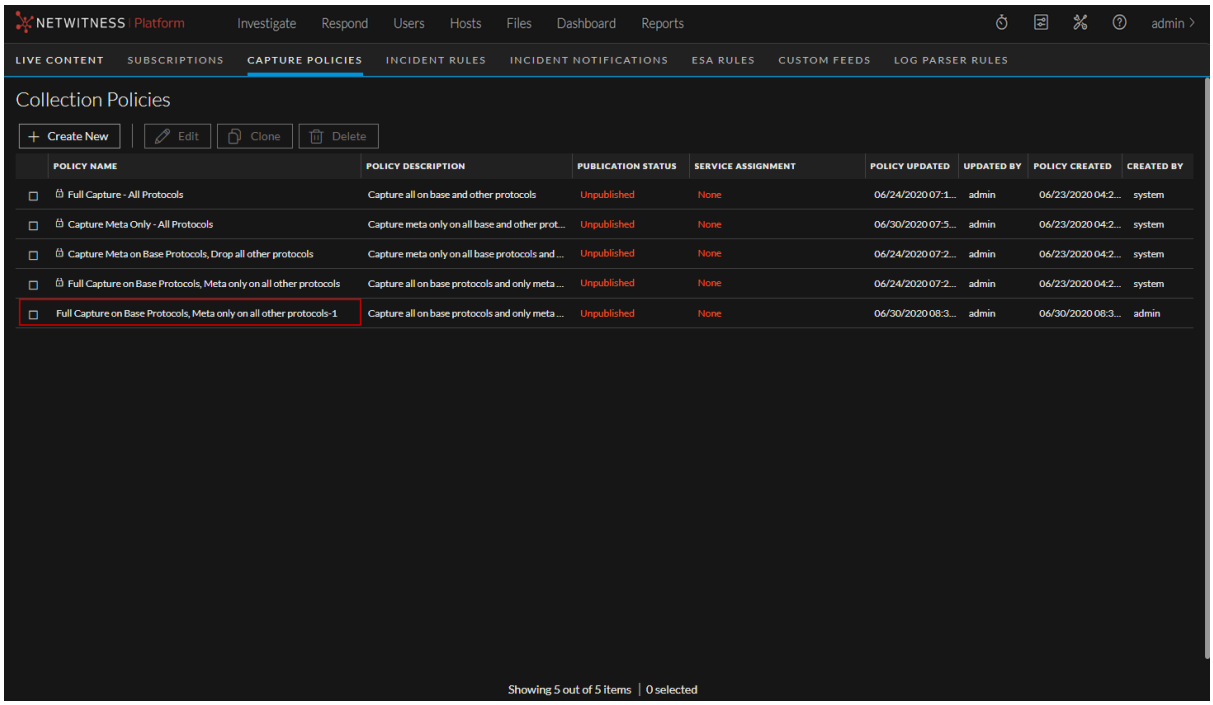
For example, as described in Use Case 2: Custom Policy for DMZ Assets , suppose a Network Decoder has been deployed into the demilitarized zone (DMZ) to protect public-facing web servers. The network communications between these web servers and other network devices are a subset of the base protocols. 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.

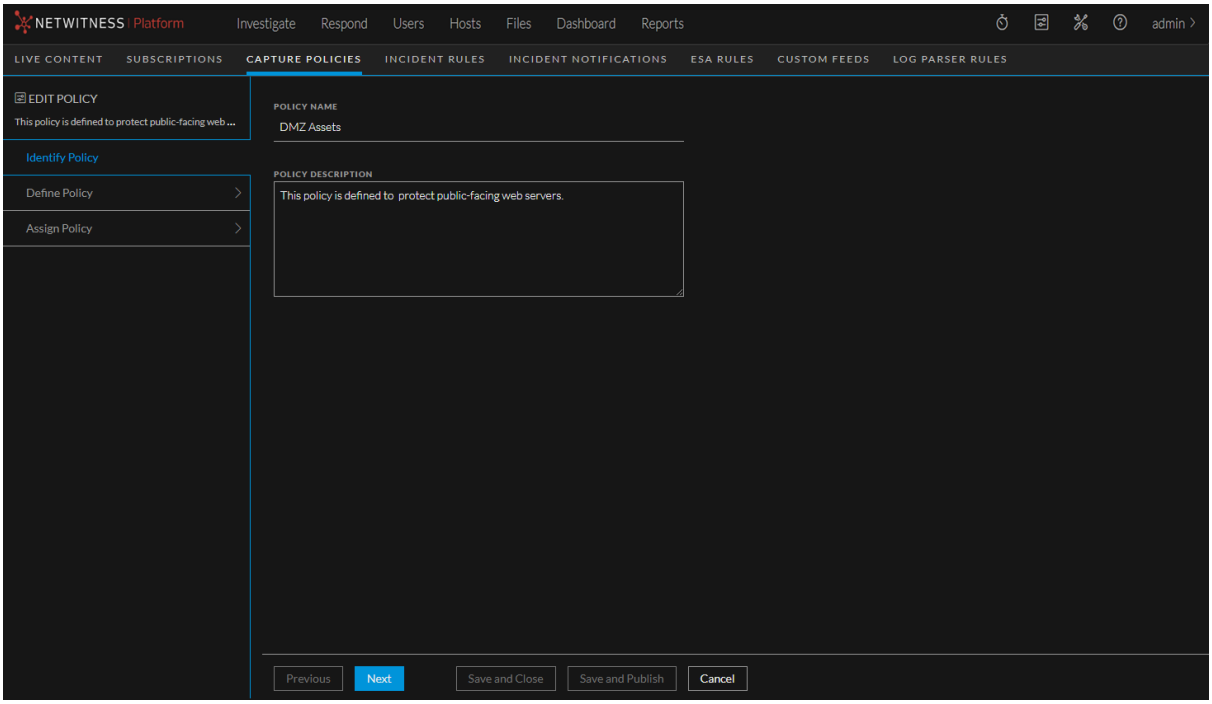
1. In the NetWitness user interface, go to (Configure) and select the **CAPTURE POLICIES** tab.



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.

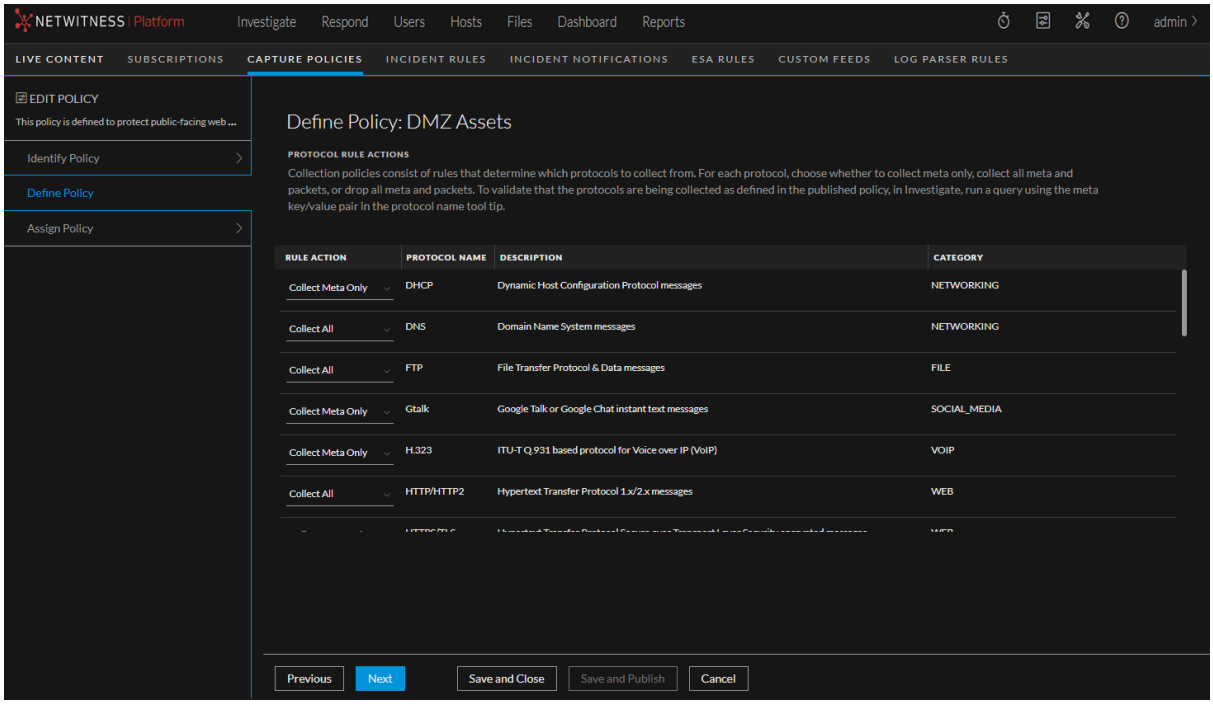


3. Select the cloned policy and click **Edit**.
   The editing view is displayed with the **Identify Policy** option active.
   In this example, we are naming this policy "DMZ Assets" and providing the description "This policy is defined to protect public-facing web servers."

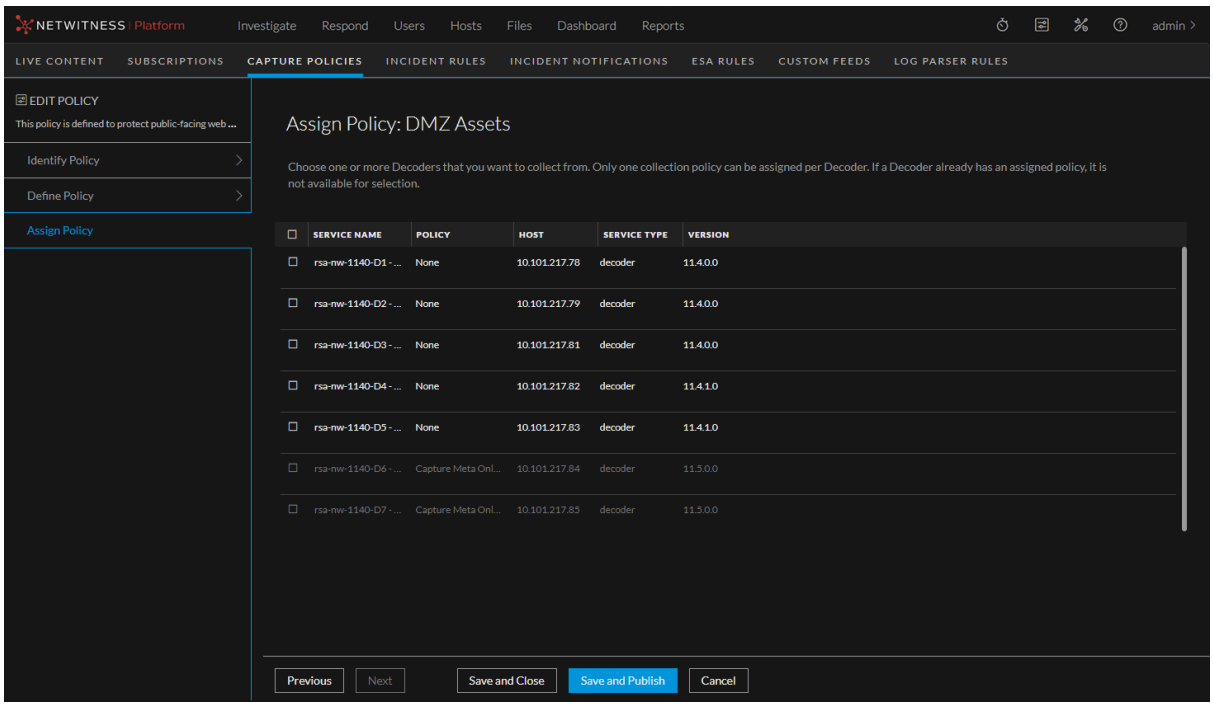


4. Click **Next**. The Define Policy view is displayed, with a list of rules you can use to define the policy. You can select the following options for rule actions:
   - **Collect Meta Only**: Collect metadata
   - **Drop All**: Drop metadata and network packets
   - **Collect All**: Collect metadata and network packets

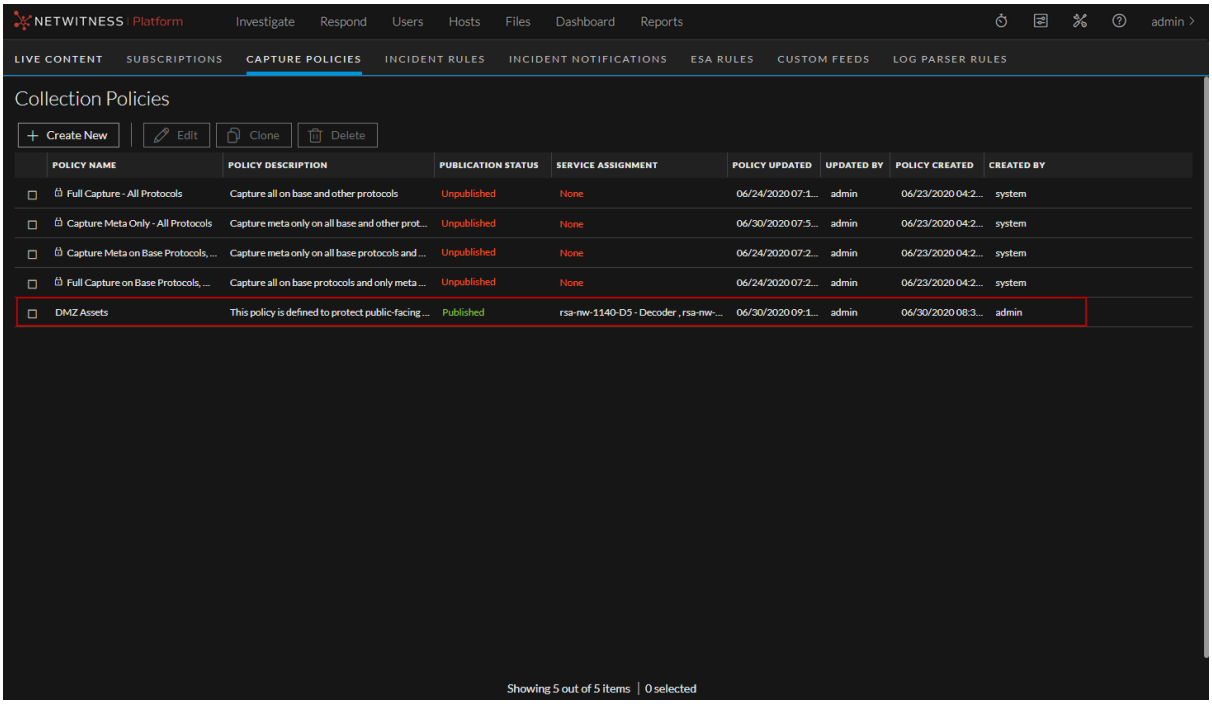
To collect the targeted data for the DMZ Assets policy, we are going to change all the base protocol rules inside the policy (except HTTP, SMTP, FTP, and DNS) from **Collect All** to **Collect Meta Only**, as shown below.



5. Click **Next**.
   The Assign Policy page is displayed.



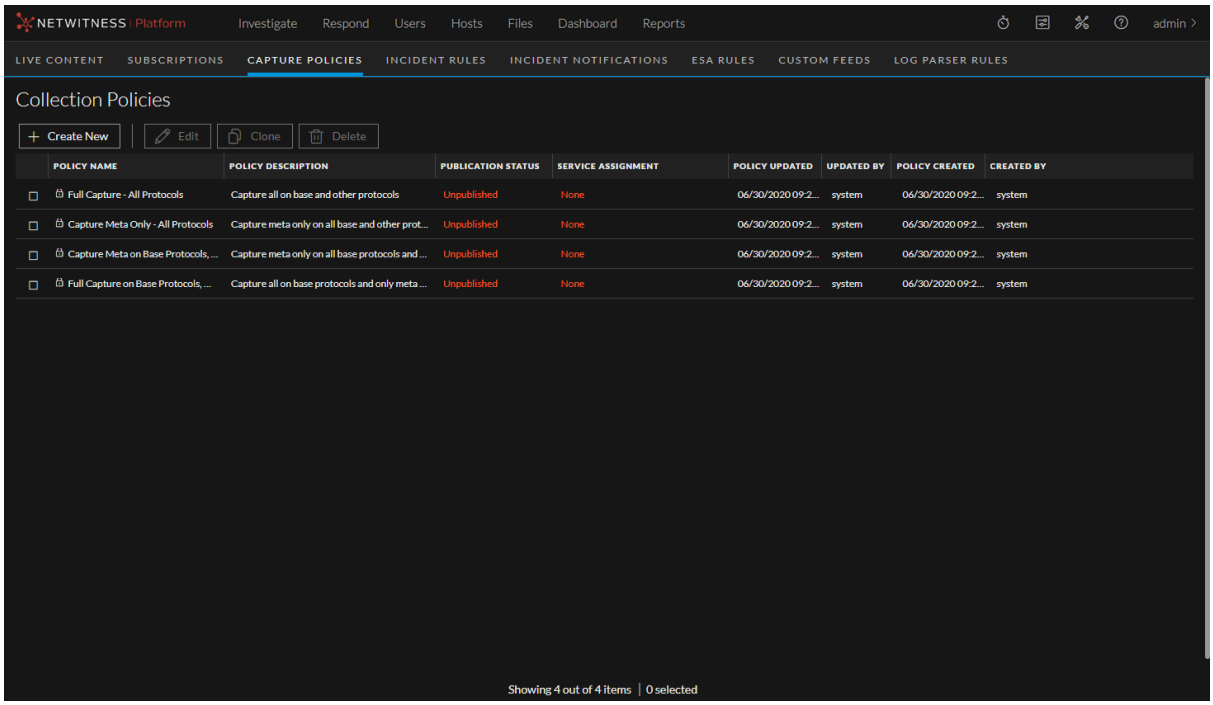
6. A Decoder can only be assigned to one policy. Decoders that are not available are greyed-out, because they are already assigned to other policies. Select the Decoders to collect the targeted data, and then click **Save and Publish**. The publication status moves to **Updating**. It changes to **Published** 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 Polices view, with the services displayed in the Service Assignment column.



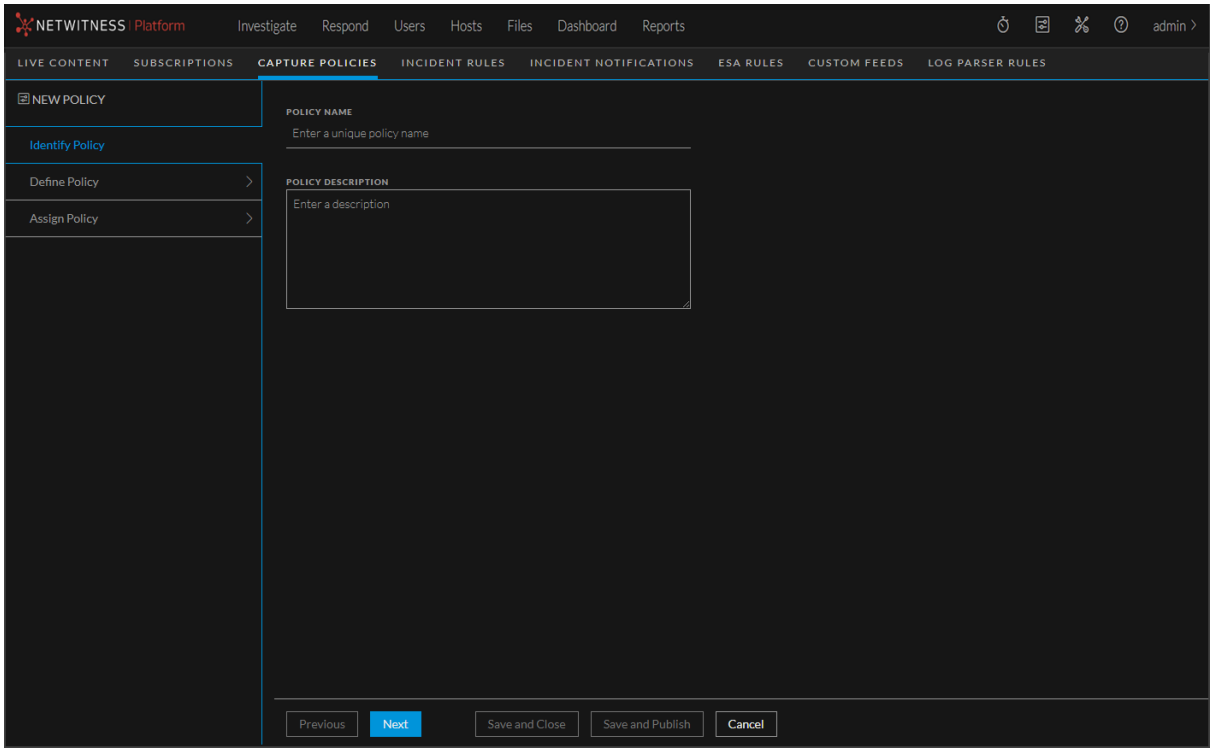
Create Custom Policies

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.

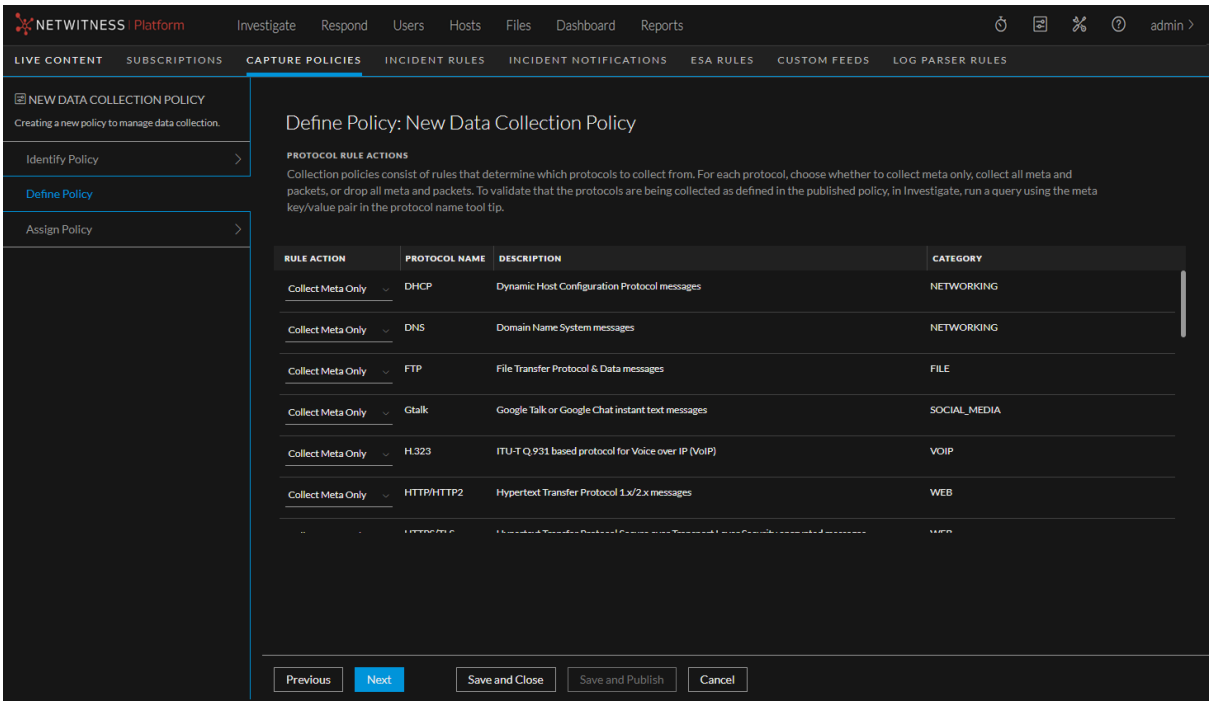
1. In the NetWitness user interface, go to (Configure) and select the **CAPTURE POLICIES** tab.



2. Click **Create New**.
   The editing view is displayed with the **Identify Policy** option active.

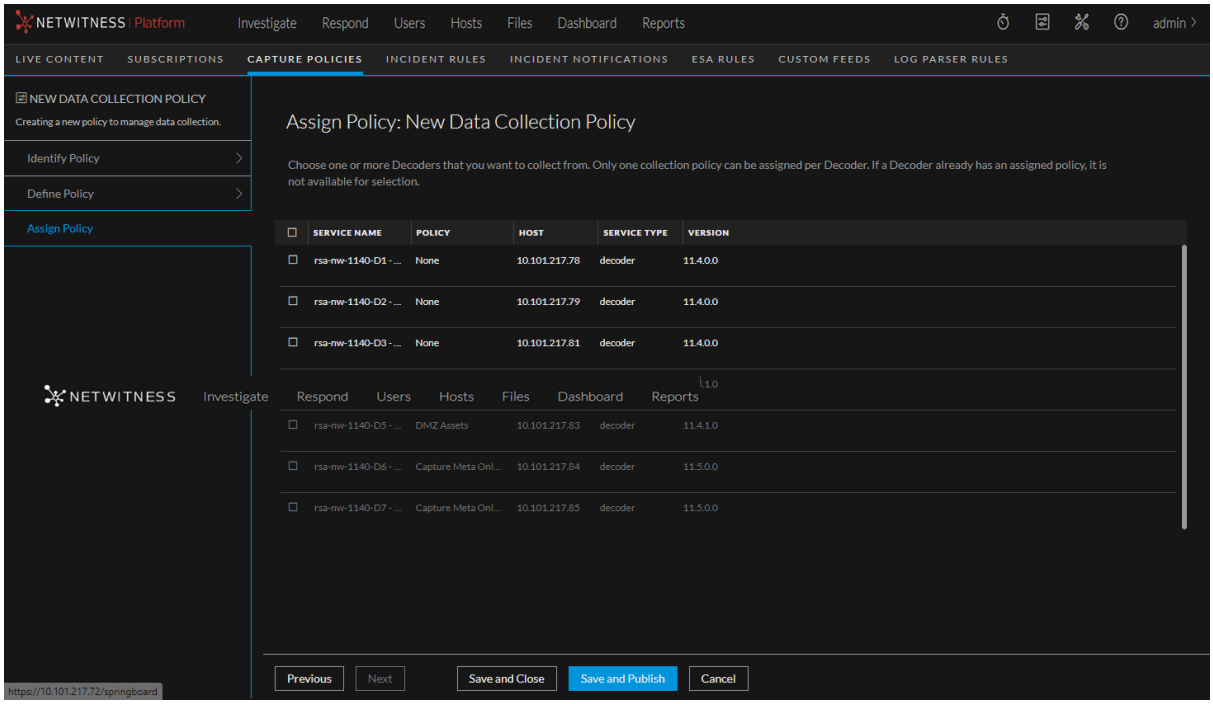


3. Enter a name and description for your policy, and then click **Next**.
   The Define Policy view is displayed.



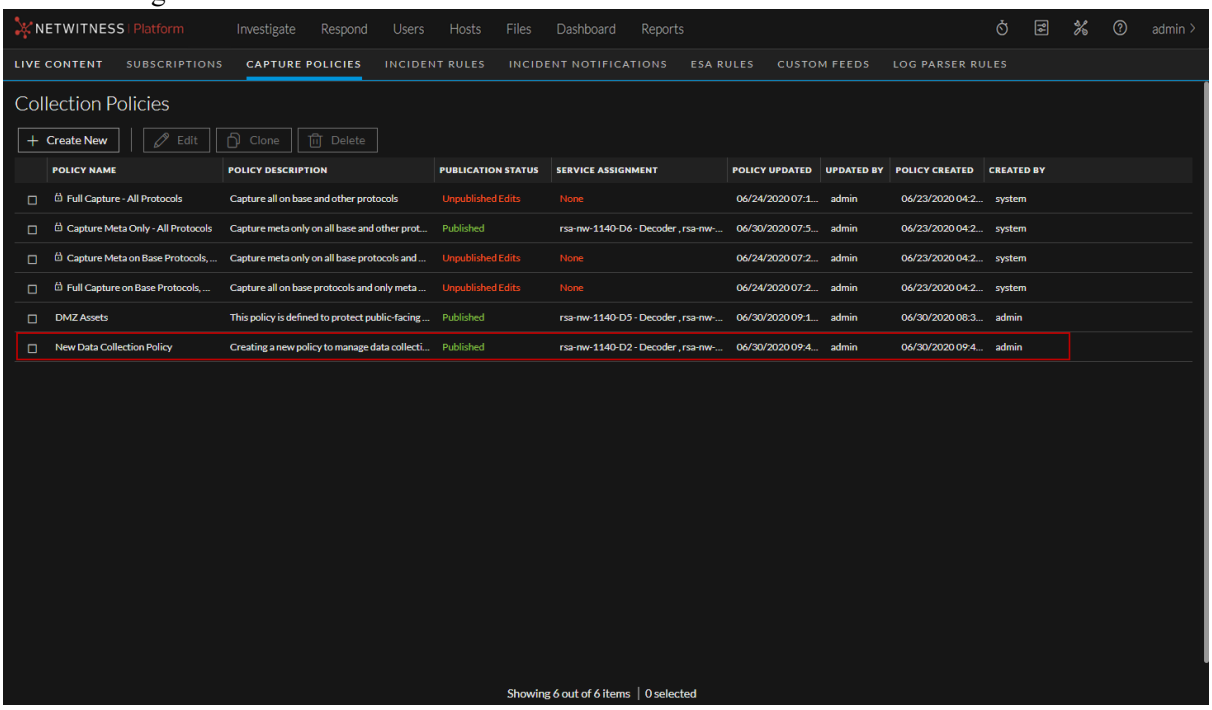
4. Define the actions for each of the protocols. The options for **RULE ACTION** are:
   - **Collect Meta Only**: Collect metadata
   - **Drop All**: Drop metadata and network packets
   - **Collect All**: Collect metadata and network packets

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 Polices view, with the services displayed in the Service Assignment column.

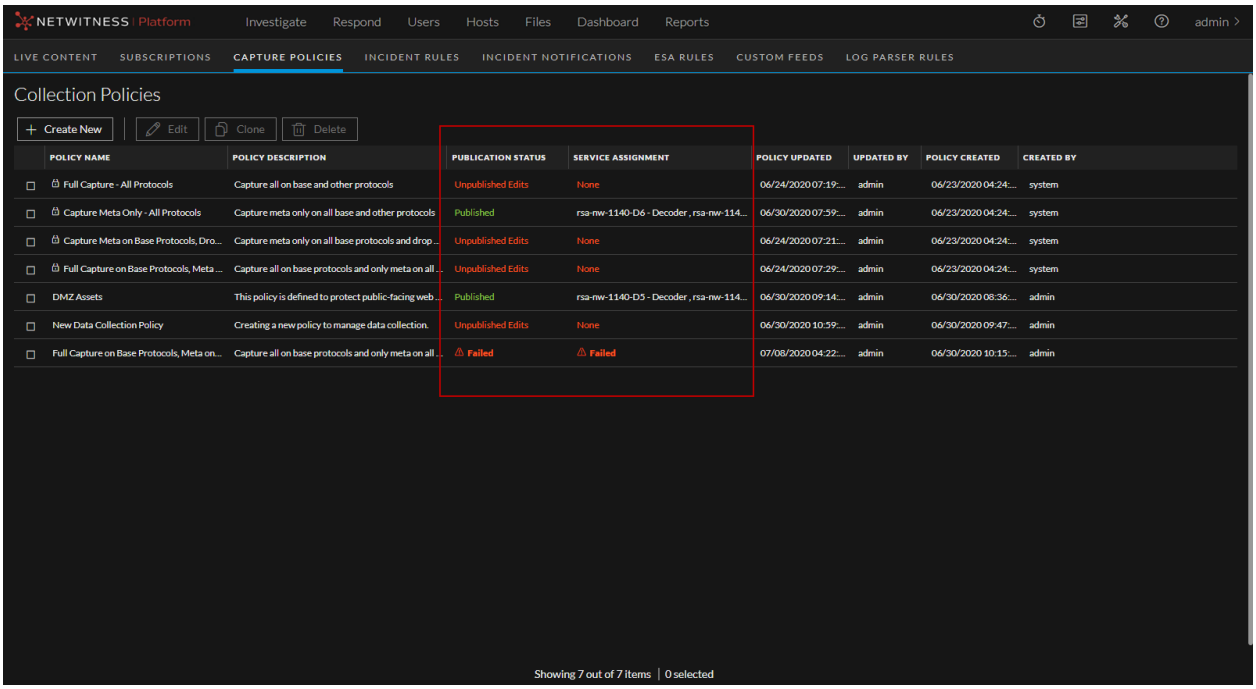


Verify Published Policies

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

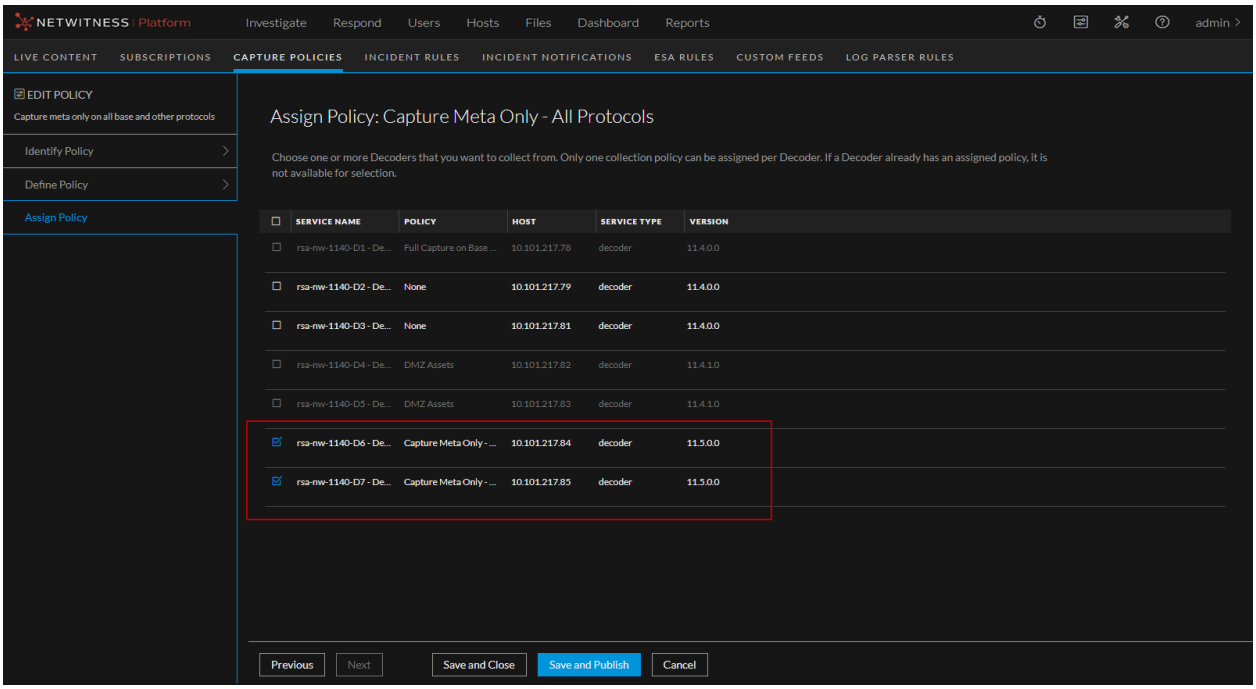
Publication status is indicated in the Capture Polices view in the PUBLICATION STATUS column. You can also see which Decoders are assigned to the policy in the SERVICE ASSIGNMENT COLUMN.



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.

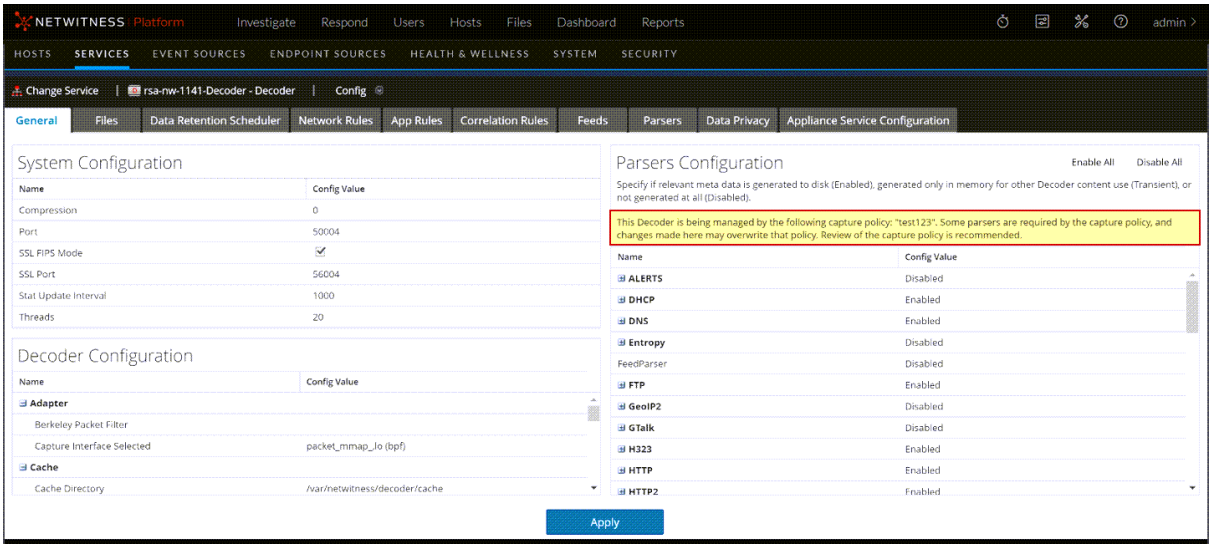
You can also select a published policy, and then click **Edit** to verify the Decoders that are assigned to the policy.



Verify Policies in Admin > Services

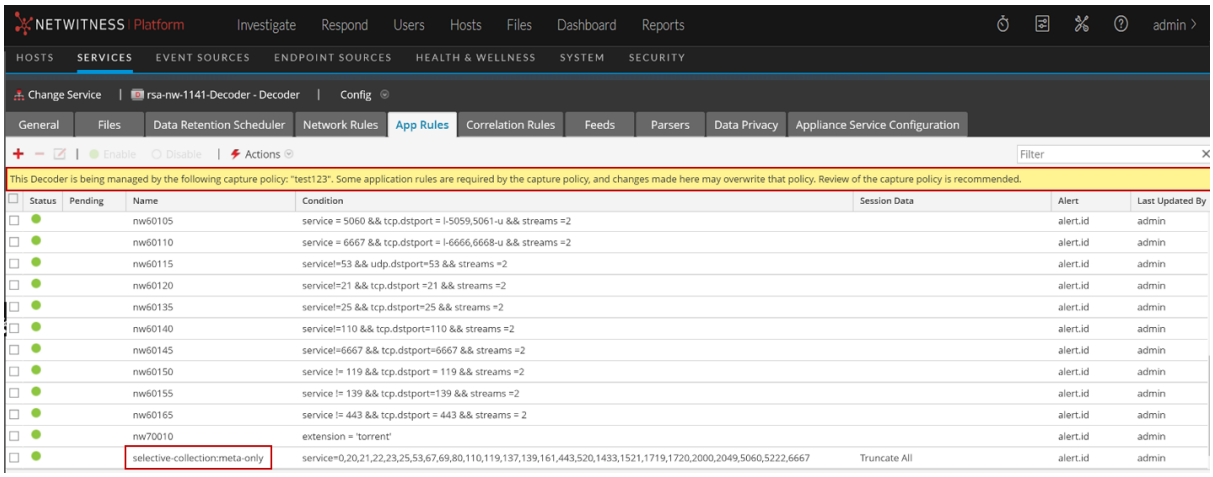
1. As an administrator, log in to the NetWitness user interface to view the Decoders that have been assigned to the policy.

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.



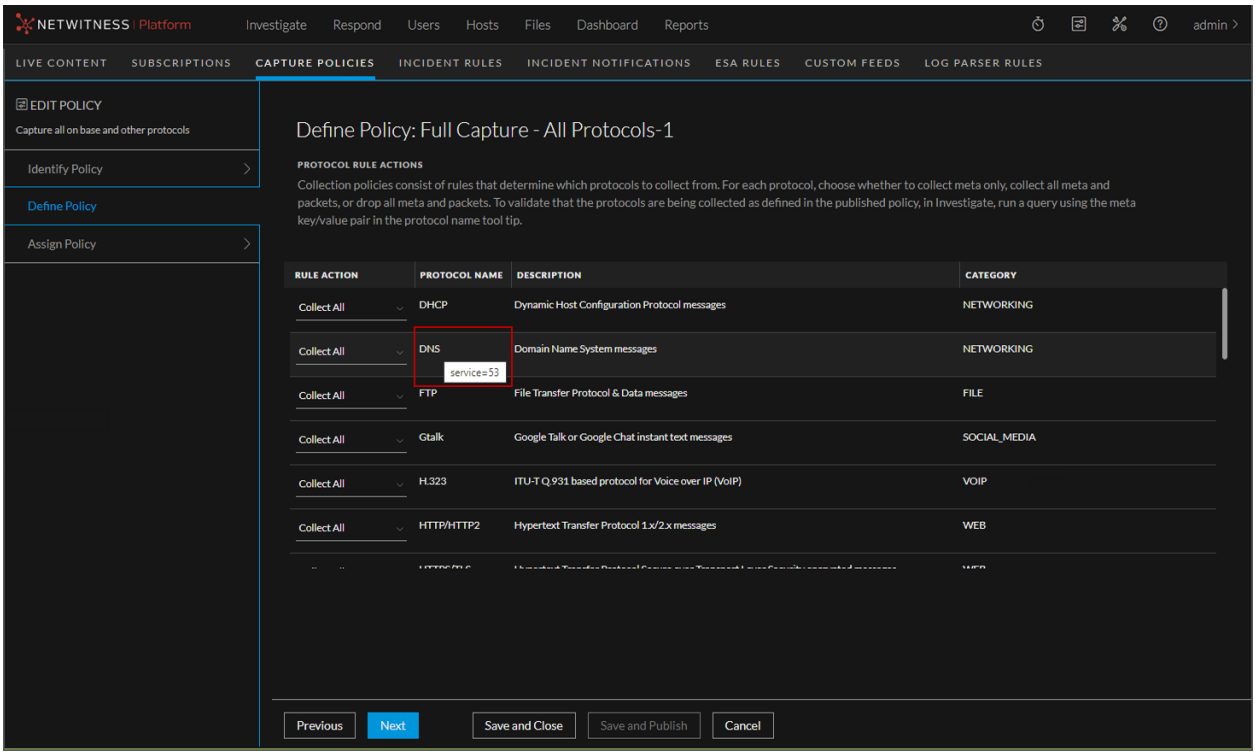
By default, the selective collection rules are added to the bottom of the rule set when the policy is applied. However, an administrator can directly affect the policy by altering the selective policy rules or by moving conflicting rules (rules with the same services) above the selective collection rules. 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.

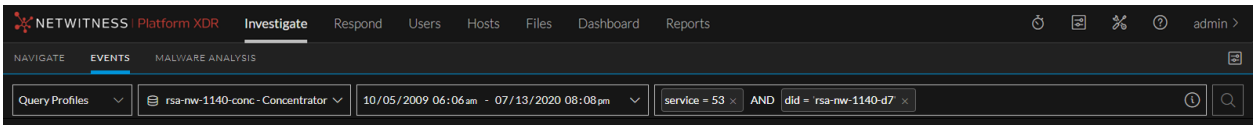
Verify Policies in the Investigate View

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.

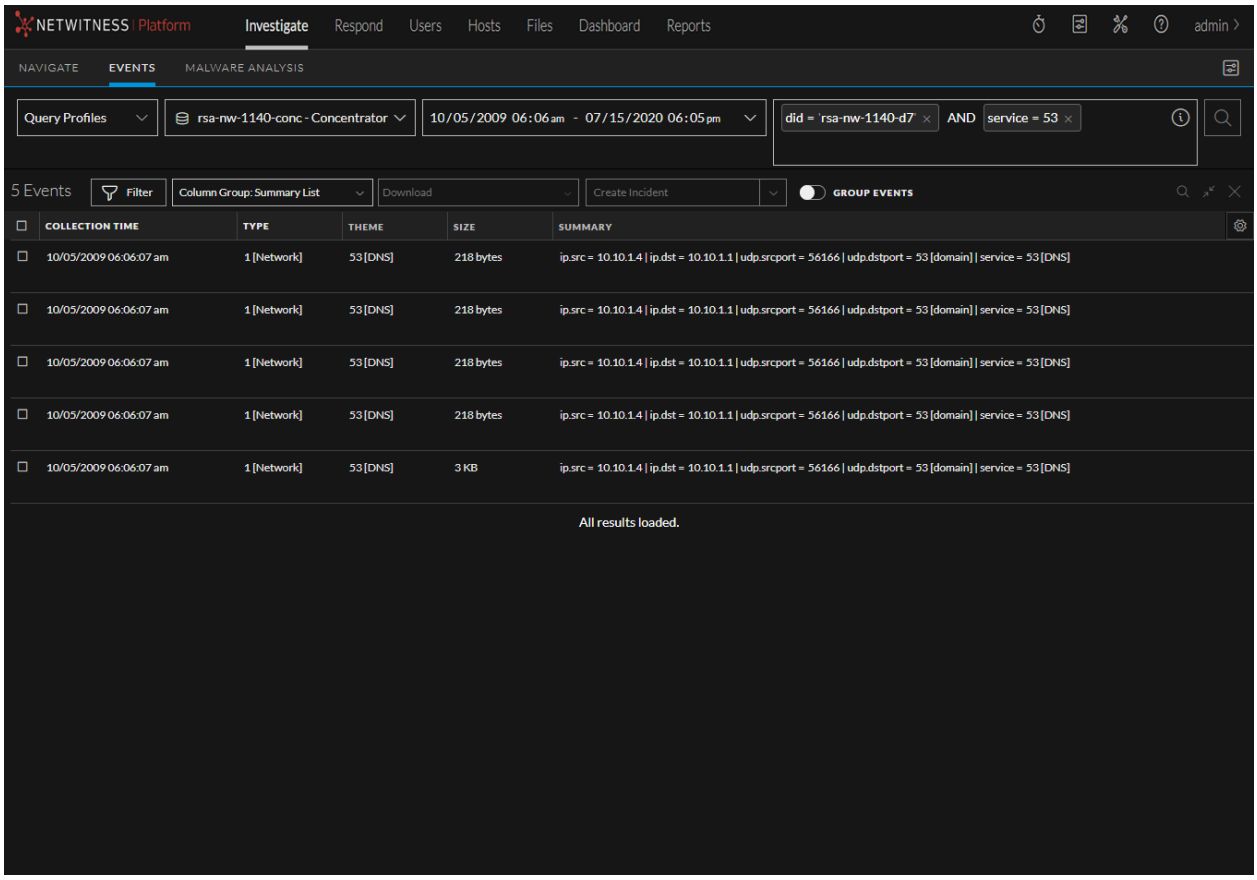
For example, in  (Configure) > **Capture Policies** > **<select a policy>** > **Edit** > **Define Policy**, if you hover over the DNS protocol name, `service=53` is displayed.



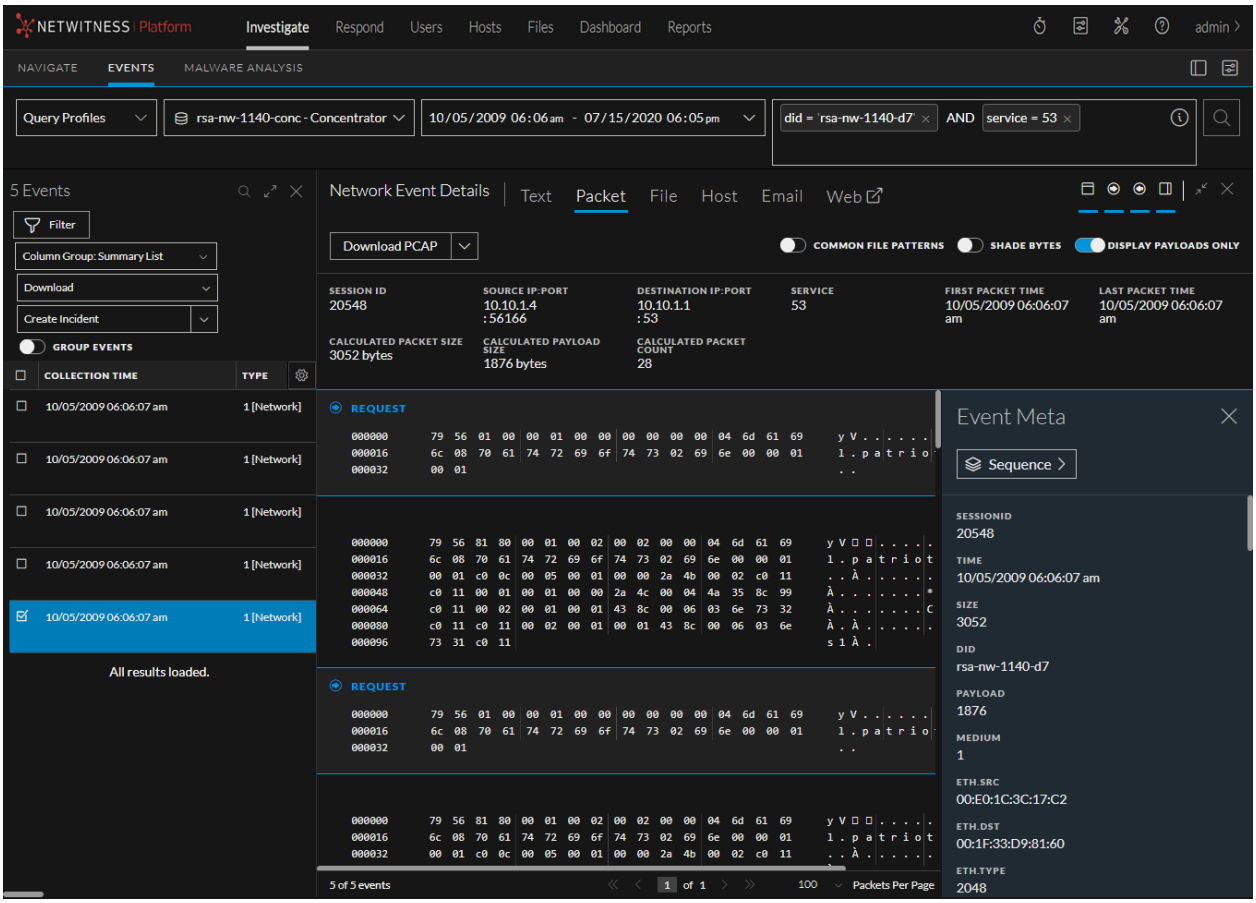
To determine if the appropriate action has been taken for DNS traffic, in **Investigate** > **Events** you can add the query parameter (for example, `service=53`) as a filter in the query bar and run the query against the Broker or Concentrator the Network Decoder is associated with for the time the 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.



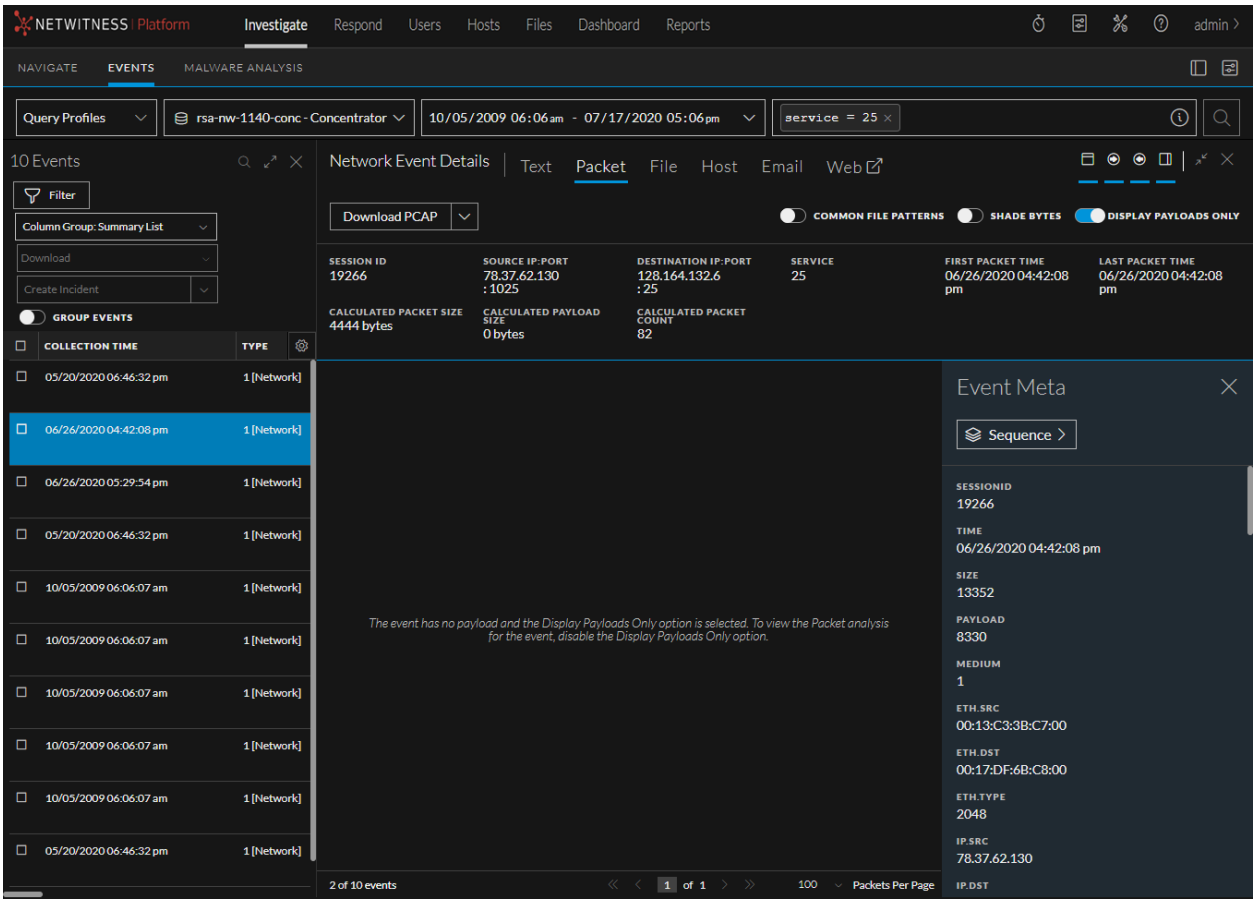
After running the query, if there are any related results, those events are displayed in the table (unless you applied an action to drop them).



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.

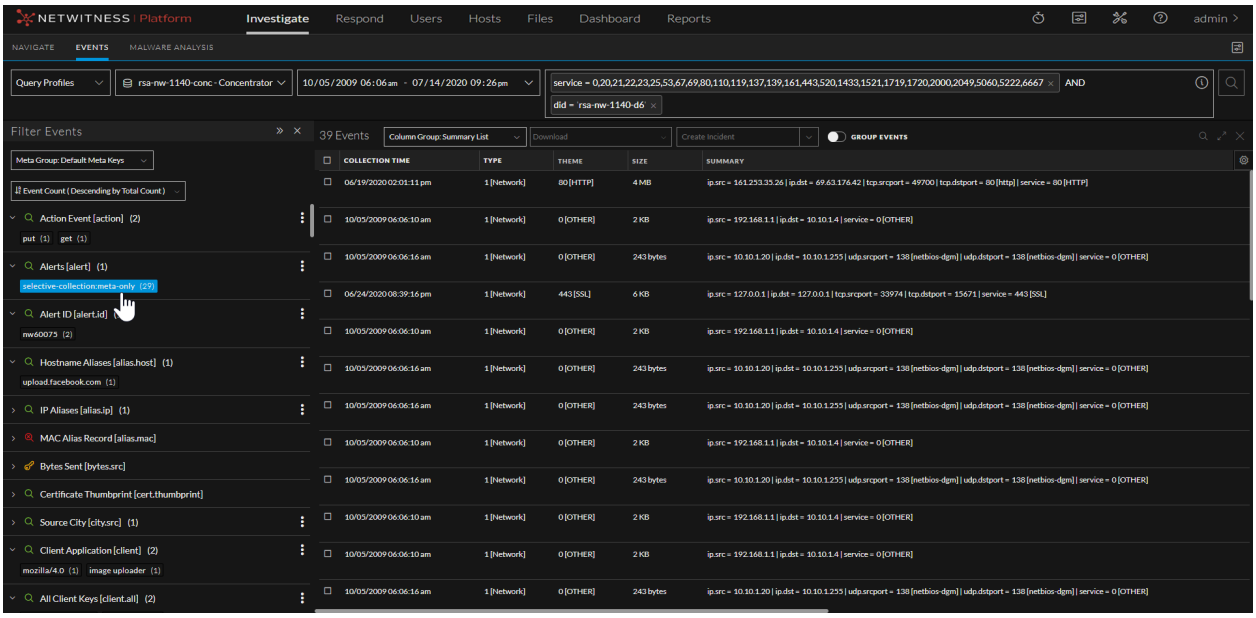


If you chose the action to `collect meta only`, you can click on one of the events and review the metadata, however, a message would be displayed instead of the packet payloads, as shown in the following figure.



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`:



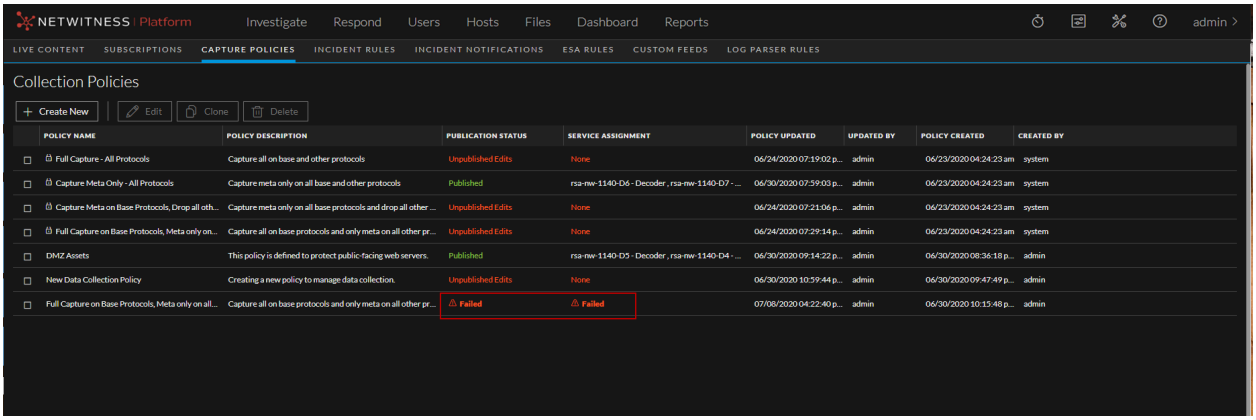
Note: If the findings do not align with what is expected from the policy, make sure that there are not other rules on the individual Decoders that might be conflicting with the selective collection policy rules. In addition, make sure when checking in Investigate that you are looking at only the traffic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

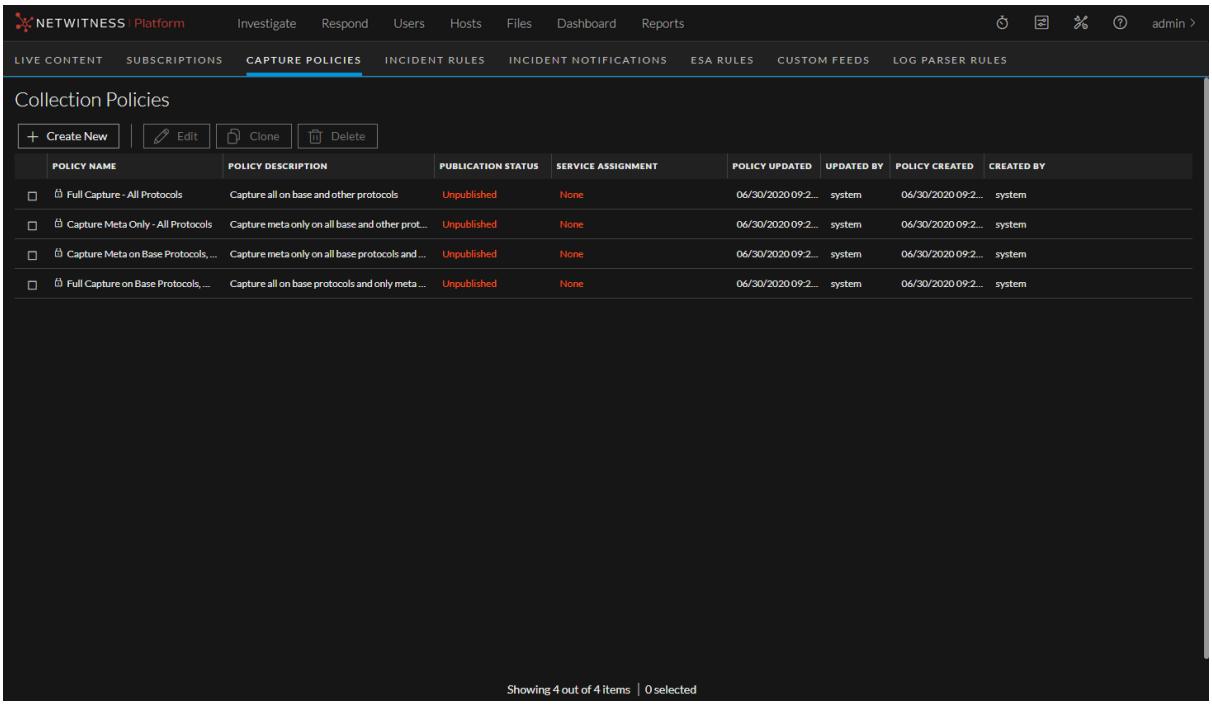
```
Jun 15 18:35:17 rsa-nw-1141-Decoder-1 NwDecoder[8260]: [Rules] [audit] User
admin (session 26182, 10.237.169.43:55380) has added application rule
'selective-collection:meta-only' at position 38
```

Unpublish Policies

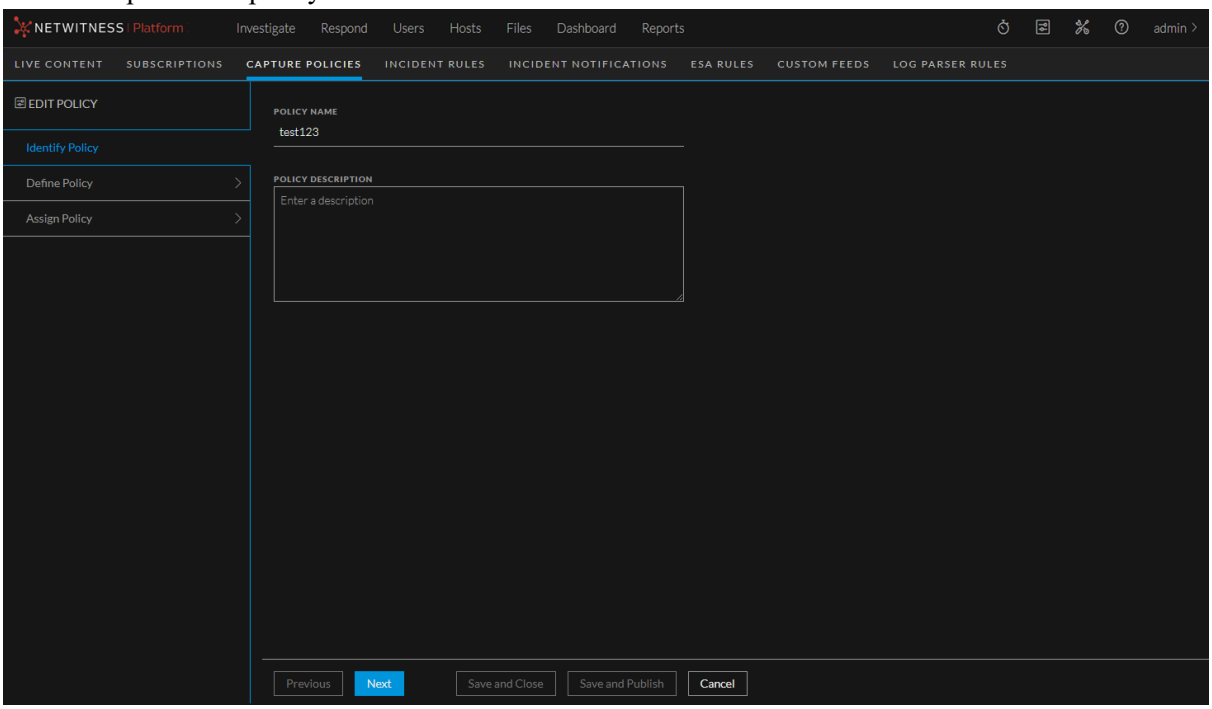
You can unpublish a policy by unassigning the Network Decoders from the policy, which preserves the policy but makes it inactive so that the data it defines is not being collected. You can then change or remove a portion of policy, or remove the entire policy from the Decoders to which they are assigned.

To unpublish a policy:

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

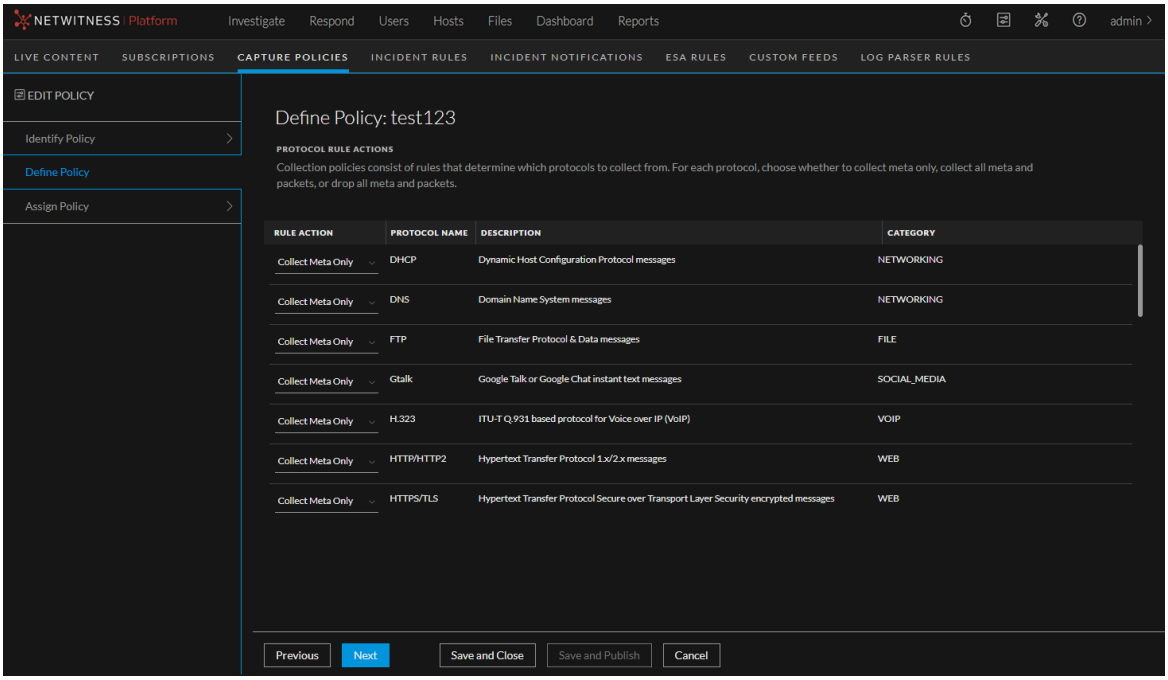


2. Select the published policy and click **Edit**.



   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**.

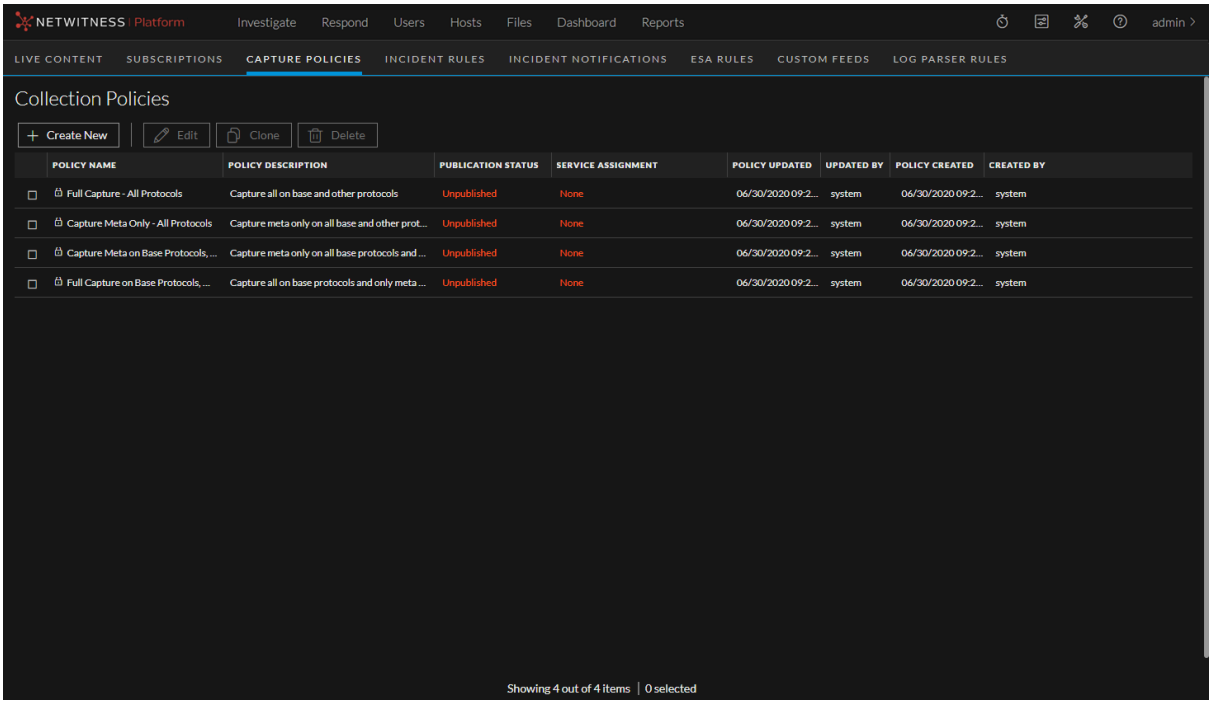
Delete Policies

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 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

Category	Protocol	Description
	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
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.

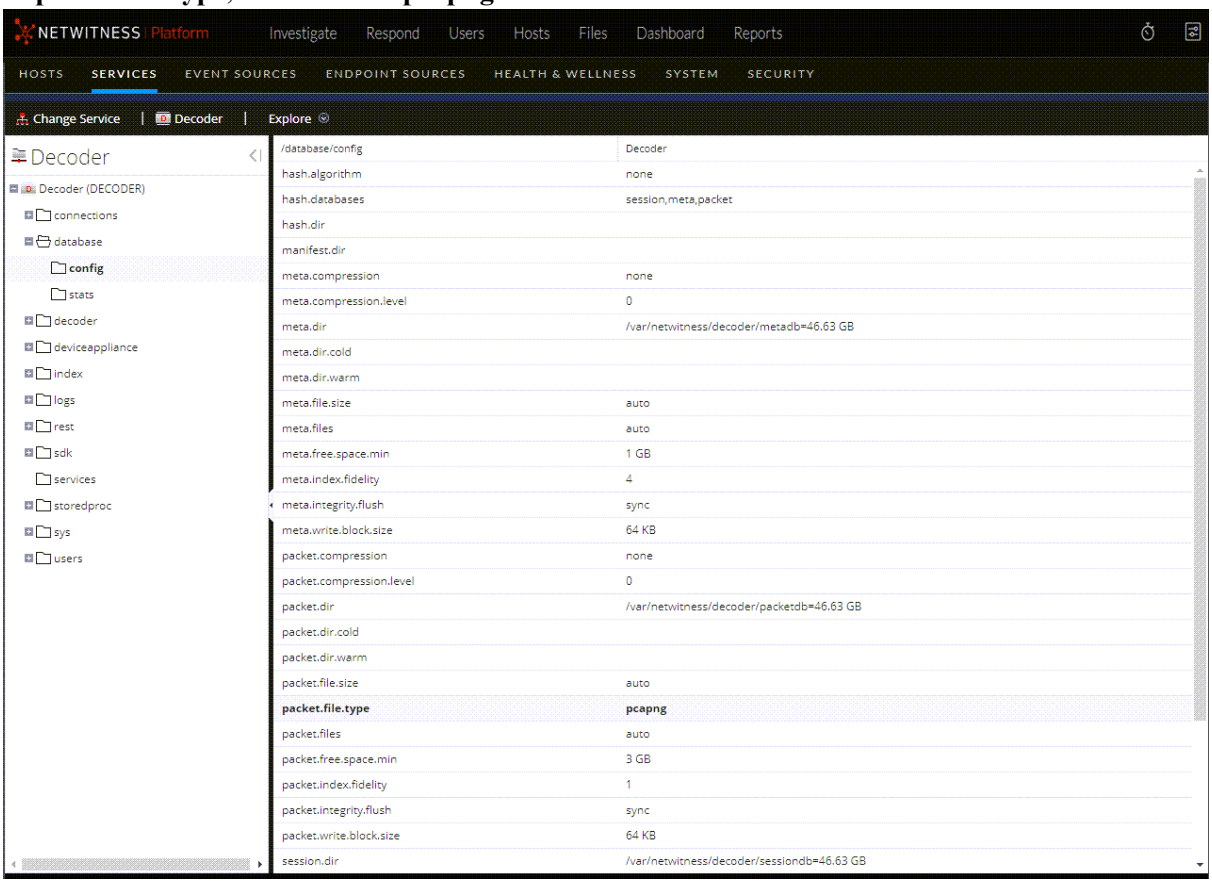
(Optional) Configure a Decoder to Write Standard pcap-formatted Files

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 latest version. If you upgrade from a previous version, 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**.



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

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, in earlier versions, 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 current version, 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.

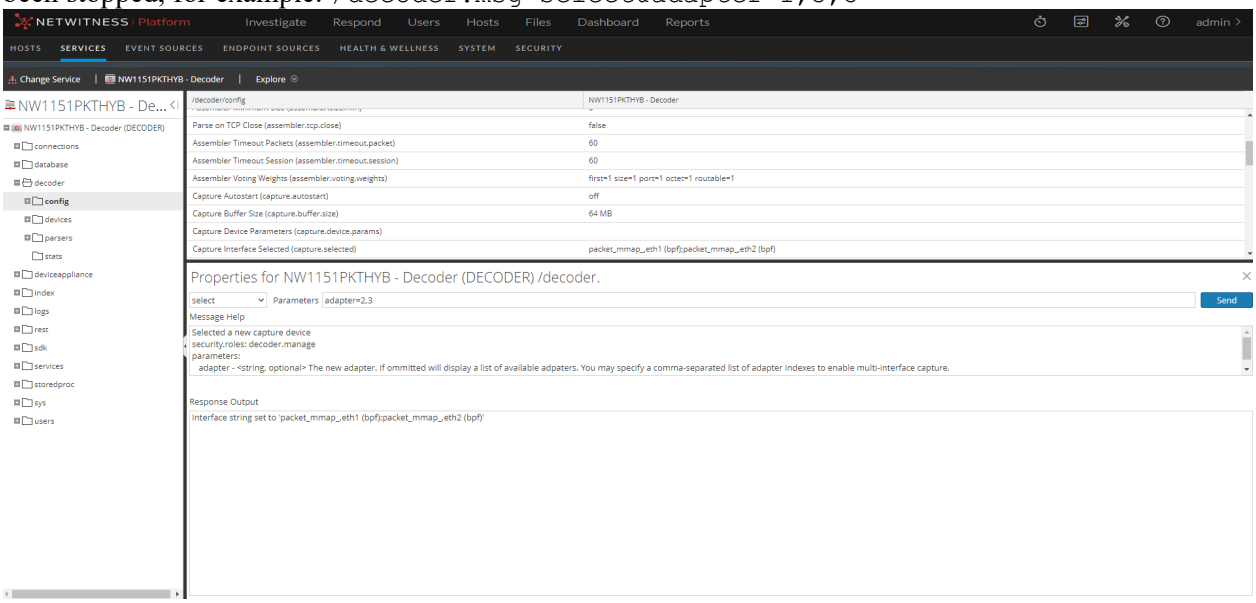
In this topic, the term "adapter" refers to a capture device within the Decoder. An adapter's name consists of the device (for example, the software API that the Decoder uses to interact with an adapter), and an interface (the physical device). A device can have multiple interfaces, and an interface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

capture.rate	SUM
capture.avg.size	MAX
capture.dropped	SUM
capture.filtered	SUM
capture.header.bytes	SUM
capture.kept	SUM
capture.payload.bytes	SUM
capture.received	SUM
capture.total.bytes	SUM
pool.packet.captured	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 NetWitness Platform*. For information about stopping and restarting capture, see Configure Capture Settings. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

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) Internet Content Adaptation Protocol Capture

Internet Content Adaptation Protocol (ICAP) is a service protocol that encapsulates HTTP messages into ICAP Messages and forwards them to an ICAP server for processing.

NwDecoder supports capturing ICAP Messages and converting the HTTP requests and responses into packets. The capture device is named `icap`, `ICAP Server`, and can be used simultaneously with other capture devices (i.e., packet capture).

ICAP Capture Options

The capture options are specified on Decoder in the configuration node `/decoder/config/capture.params`

You must restart the capture before any changes take effect.

- `reqmod=<bool>` - By default, the Decoder will process both `REQMOD` and `RESPMOD` messages. Since the `RESPMOD` contains both the request and response, it may be advantageous for Decoder to ignore `REQMOD` messages and only process `RESPMOD`. If you set the `reqmod=false`, the Decoder will generate a request packet and response packets from a single `RESPMOD` message. The default behavior is equivalent to `reqmod=true`

- `client_ip=<string>` - You can configure some ICAP Clients to include the originating IP address in the ICAP headers. The name of the header with this information can vary, so you can use this setting to specify the name of the client IP header. For example, `client_ip=X-Client-IP` would configure the Decoder to extract the IP from the `X-Client-IP` header. Ignoring this setting will disable searching for the header.

(Optional) Data Plane Development Kit Packet Capture

The NetWitness 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, 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 in earlier versions, 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**.

4. 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.



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

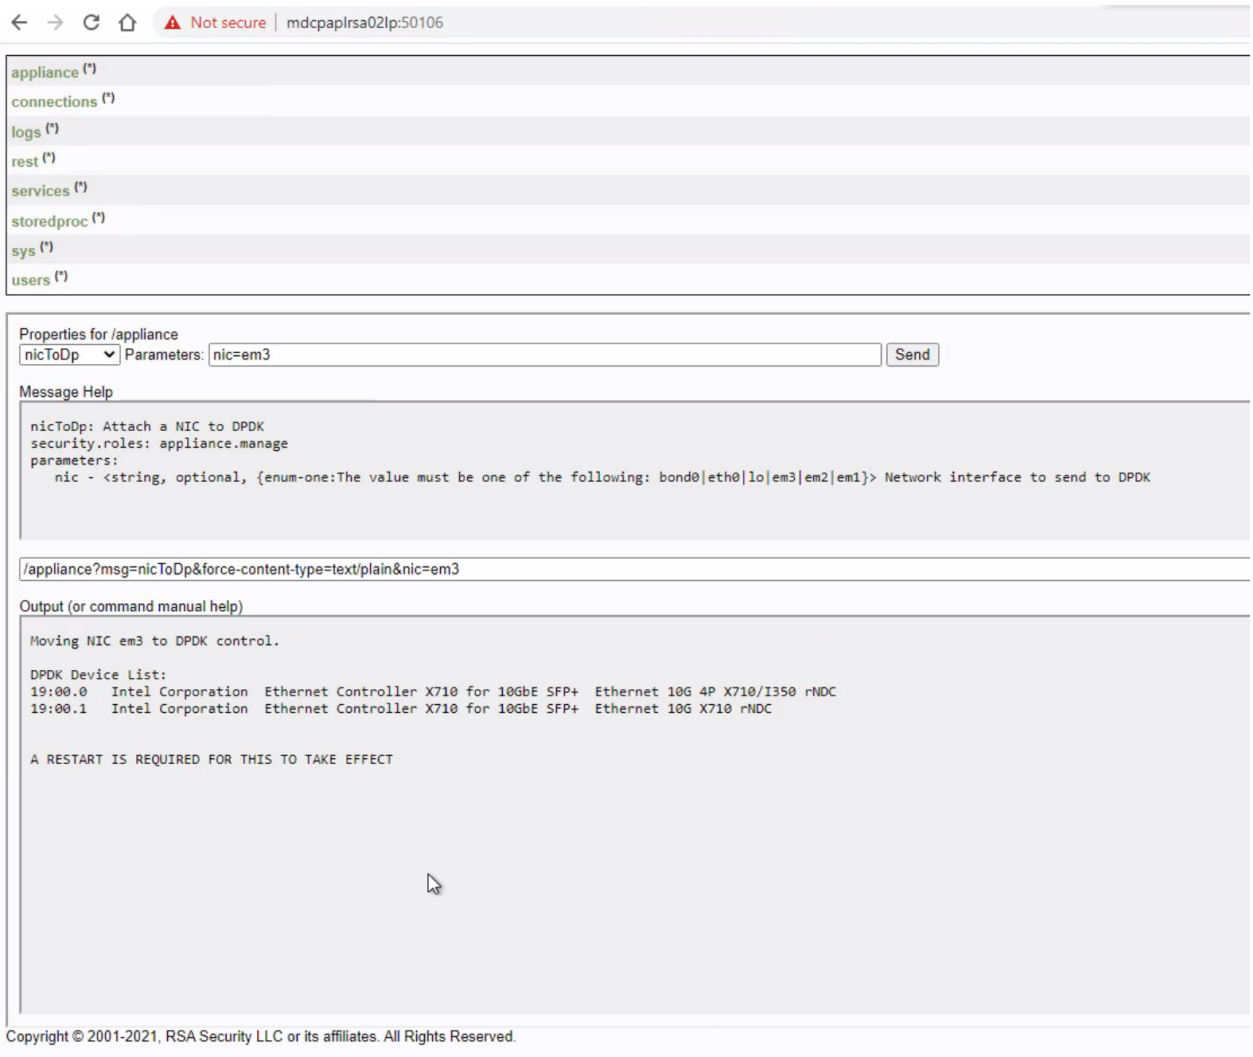
> **Note:** The `prune` command removes only the undetected interface folders. Also, it ignores the `interfaces`, `stats`, `nwimport`, and `flow_events` folders available in `/decoder/devices/`.

Manually Move a Capture Interface to DPDK

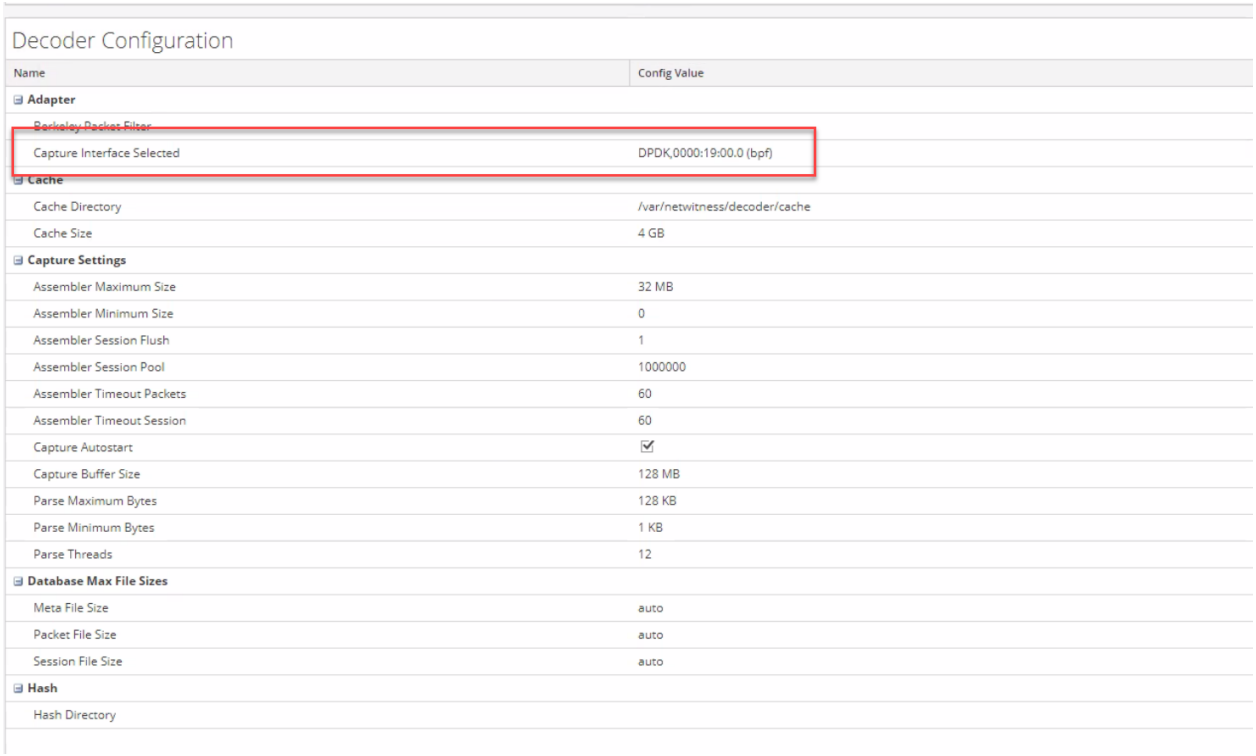
1. Run the new host service command `/appliance/nicToDp`, that converts a named network interface and sends it to DPDK.
   For example, if you provide `nicToDp` an interface name like `p1p1` it detaches it and sets it up for DPDK to use. It does the following:
   a. Identifies the PCI address of the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



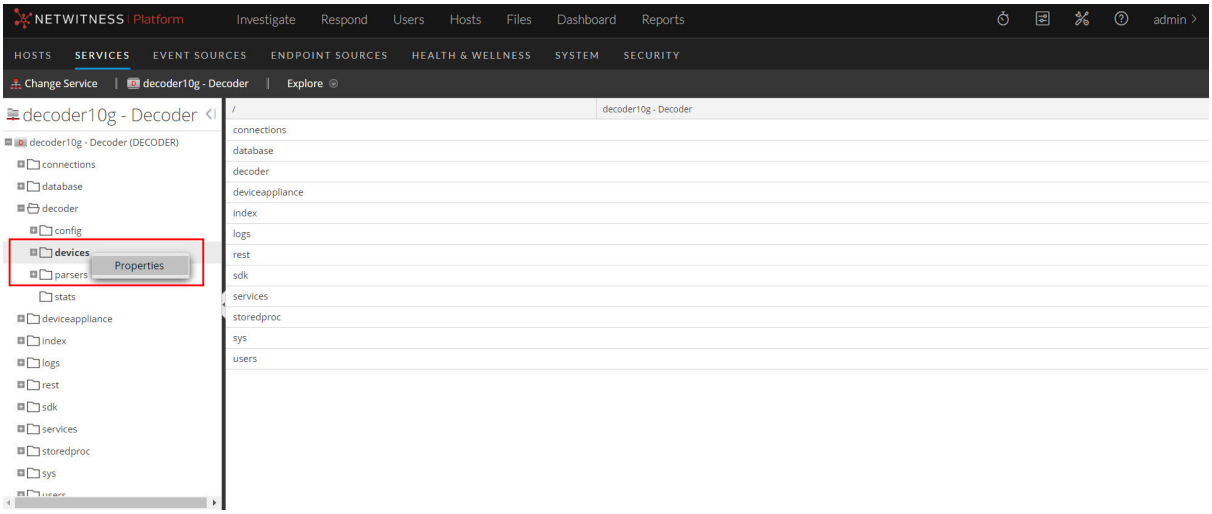
2. Reboot the Decoder host. The reboot accomplishes the following:
   a. Ensures that live capture on the interfaces stops before the network driver is unloaded.
   b. Binds the `vfio_pci_generic` driver to the configured interfaces.
   c. Pre-allocates some Huge Page memory for each interface, as required 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. The following figure shows the Decoder configuration after a restart.



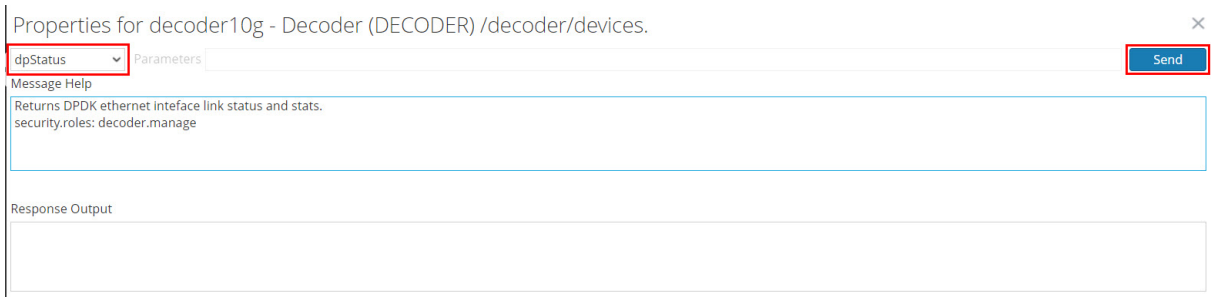
Check the device status and stats of DPDK interfaces

When you capture using DPDK interfaces, the link status and stats of the DPDK can be obtained using the command `/decoder/devices dpStatus`. The traditional commands like `ethtool` or `ifconfig`, which are used for other capture interfaces, will not fetch the link status and stats of DPDK devices.

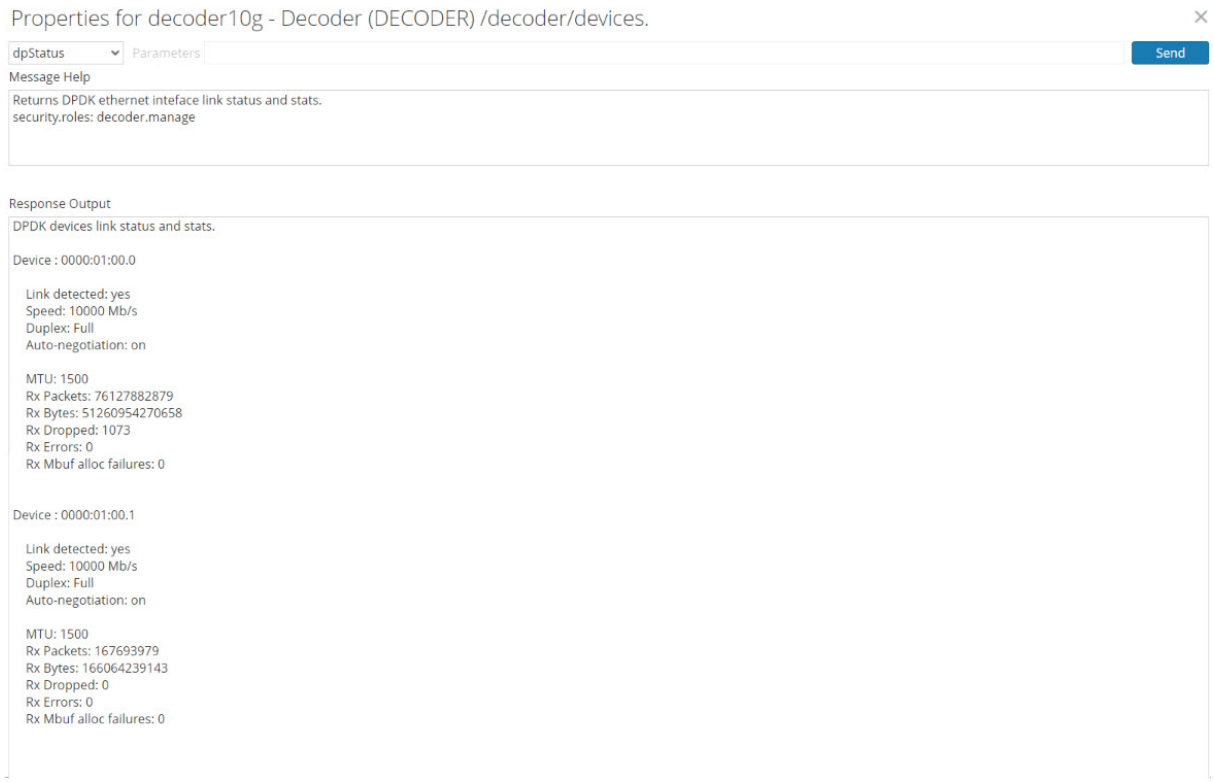
1. Go to (Admin) > **Services**.
2. In the Administration Services view, select the Decoder and **View > Explore**.
3. Navigate to **decoder > devices**, right click on **devices** and select **Properties**.



4. Select **dpStatus** from the drop-down and click **Send**.



The link status and stats of DPDK devices are displayed in the **Response Output** section.



Check link status and stats from NwConsole

Run `/decoder/devices dpStatus` on NwConsole. Sample output when capture is in **Start** state:

```
[localhost:56004] /> send /decoder/devices dpStatus
DPDK devices link status and stats.

Device : 0000:01:00.0

    Link detected: yes
    Speed: 10000 Mb/s
    Duplex: Full
    Auto-negotiation: on

    MTU: 1500
    Rx Packets: 416843663
    Rx Bytes: 281110914127
    Rx Dropped: 314
    Rx Errors: 0
    Rx Mbuf alloc failures: 0


Device : 0000:01:00.1

    Link detected: yes
    Speed: 10000 Mb/s
    Duplex: Full
    Auto-negotiation: on

    MTU: 1500
    Rx Packets: 597764
    Rx Bytes: 373053232
    Rx Dropped: 0
    Rx Errors: 0
    Rx Mbuf alloc failures: 0
```

Sample ouput when the capture is in **Stop** state:

```
Device : 0000:01:00.0
Link detected: no
Device : 0000:01:00.1
Link detected: no
```

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`.

Configure the Decoder Capture Failover in Load Balance Deployments

The DPDK capture in the decoder service provides the ability for load balancers to failover and avoid data loss when decoder capture is off or when the service is shut down due to maintenance activity. When the decoder capture is stopped, the DPDK capture interface link status is down. This provides the ability for the traffic ingestion physical load balancers (for example, Keysight Ixia, Gigamon) which monitor interface links to switch to a different decoder and avoid data loss.

Configuration Failover on DPDK Capture Interfaces

Refer to the (Optional) Data Plane Development Kit Packet Capture *topic for instructions to configure DPDK capture on the decoder*.

- You can migrate existing PF_RING devices to DPDK

  (Or)

- Manually move capture interfaces to DPDK

Disable Wake On LAN (WOL) in BIOS settings

Some load balancers monitor interface link status through NIC link lights. You must disable WOL as the interface will not be down if WOL is enabled. To disable WOL in the BIOS, follow the BIOS instructions.

Example: f2 -> Device Settings -> "Integrated NIC 1" (capture interfaces) -> "NIC Configuration" -> "Wake on LAN" = Disabled

Note: When you restart the device and if capture does not auto-start, the lights on the capture interface are on. But if you turn on capture and turn it off again the lights will be off.

(Optional) Remove DPDK on Capture Interfaces

In case you need to revert DPDK on capture interfaces you can use the Appliance service `dpToNic` command.

- If you run this command without arguments, it lists the ethernet devices controlled by DPDK.
  - **Example:** /appliance dpToNic

    DPDK Device List:

    82:00.1 Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection Ethernet Server Adapter X520-2
- If you specify an ethernet interface using the 'eth' parameter, it removes the interface from DPDK so that when the system reboots it will be treated as a normal Linux network interface.
  - Example: /appliance dpToNic eth=82:00.1

    Detaching PCI Device 82:00.1 from DPDK.

    DPDK Device List: A RESTART IS REQUIRED FOR THIS TO TAKE EFFECT

Validate Link Status

- On Capture Start
  - During the decoder capture start, the service logs DPDK interface link status
    - [DPDK] eth dev 0000:05:00.0 link is up
  - On the upstream traffic ingest Load Balancer, the link status would be detected.
    - If it is a Linux-based device, you can use the ethtool <interface name> command to check link status
- On Capture Stop
  - During the decoder capture stop, the service logs that capture is stopped on DPDK interfaces
    - [Thread] Stopped thread: Capture Work (DPDK, 0000:05:00.0) id: 31214
  - On the upstream traffic ingest Load Balancer, the link status would not be detected.
    - If it is a Linux-based device, you can use the ethtool <interface name> command to check link status.

DPDK Interface Configuration and Statistics

- The `dpdk-devbind --status` command displays the status of DPDK and other known network interfaces. It displays the PCI domain, bus, slot, and function, along with a text description of the device.
  - Example: `dpdk-devbind --status`
    - 0000:01:00.0 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' drv=vfio-pci unused=ixgbe_zc
    - 0000:01:00.1 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' drv=vfio-pci unused=ixgbe_zc
- The default MTU is 2048. If you need to adjust the MTU for capture, add the `snaplen` parameter to `/decoder/config/capture.device.params` .
- VLAN offload is disabled.
- The DPDK interface capture stats are updated under /decoder/devices/dpdk/
  - Example: /decoder/devices/DPDK/0000:01:00.1/stats/
    - Received packets (Rx packets) - capture.received
    - Received bytes (Rx bytes) - capture.total.bytes
    - Received dropped (Rx dropped) - capture.dropped

For information on checking device status and stats of DPDK interfaces, refer to Check the device status and stats of DPDK interfaces

(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.

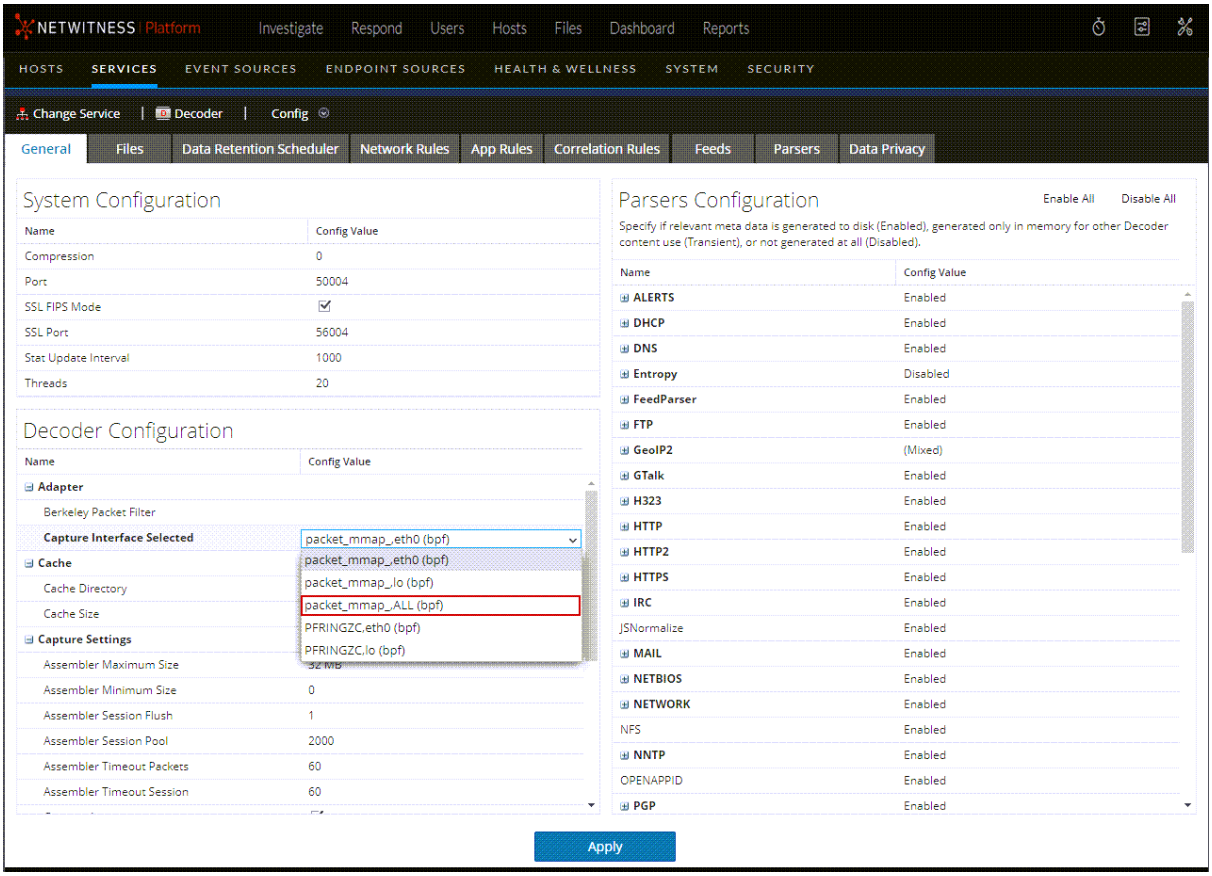
These are `vlan` examples of both options. If you need to pass multiple settings for `capture.device.params`, use the following synax. Notice that quotes are needed to delineate whitespace. For more information about how to use quotes for whitespace, see the "Connecting to a Service" topic in the *NwConsole User Guide for NetWitness Platform*.
`name1="value1" name2="value2"`.

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

To configure the `packet_mmap_` adapter to preserve the VLAN tags in packets:

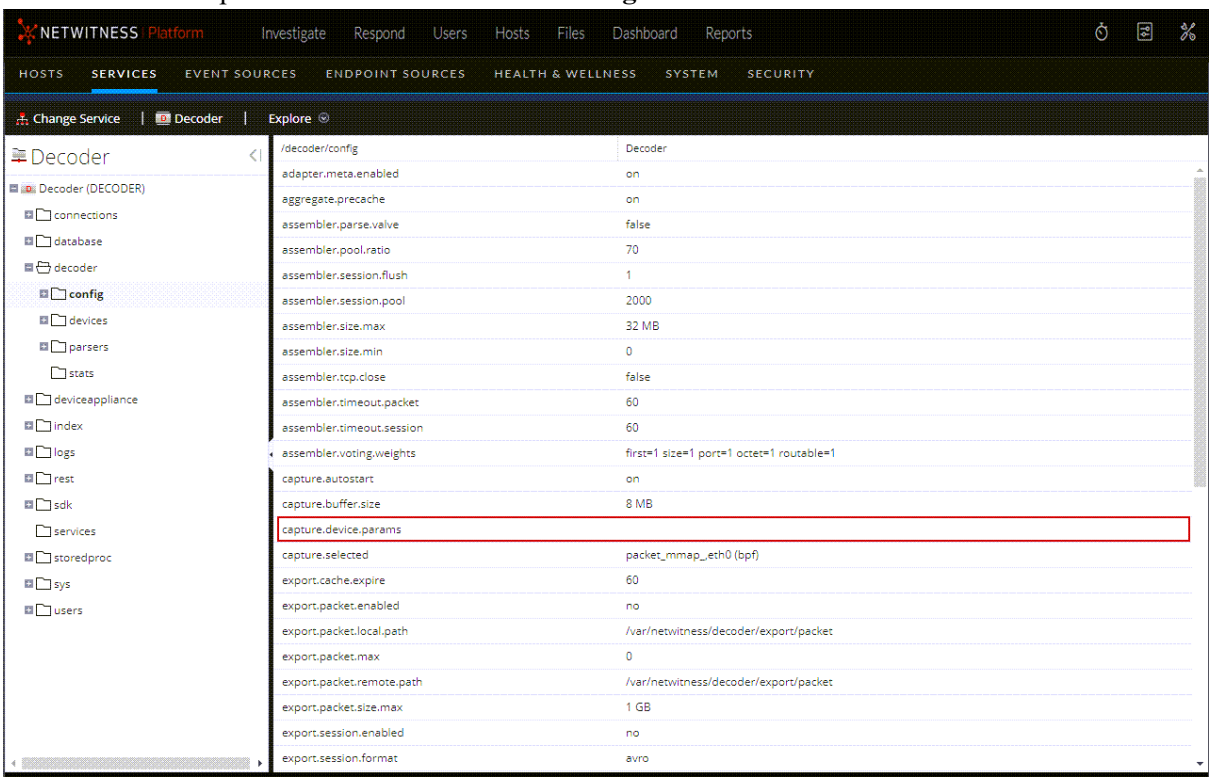
1. Go to  (Admin) , select the Decoder service and  **> View > Config.**

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



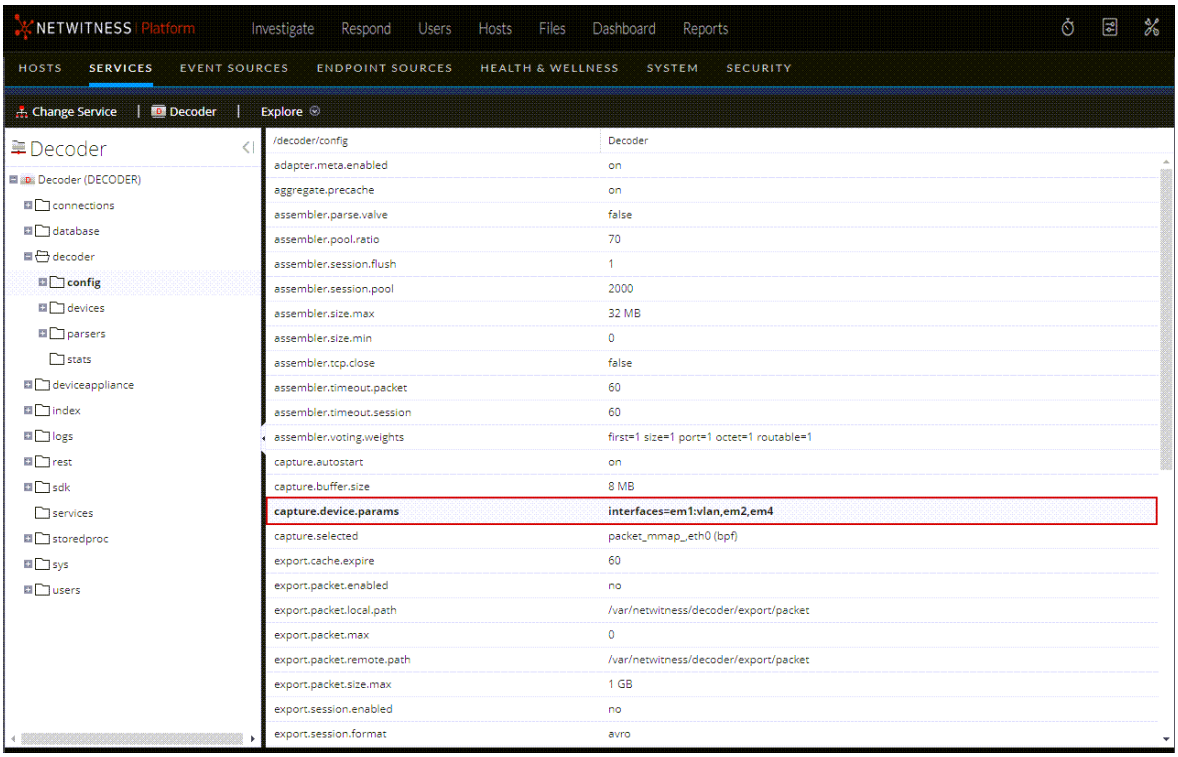
3. To go to the **Explore** view, click **Config** in the toolbar and select **Explore** in the drop-down list.

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



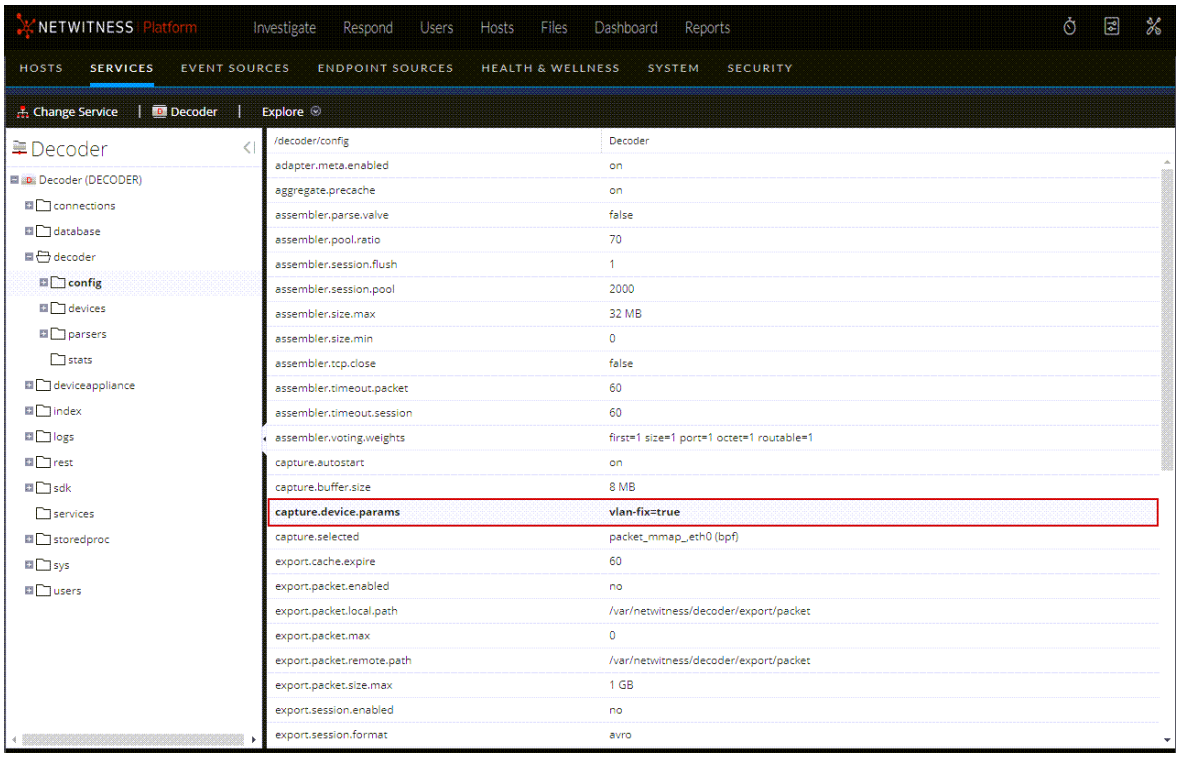
5. Click in the values column next to `capture.device.params`, and do one of the following:

   - 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`



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.

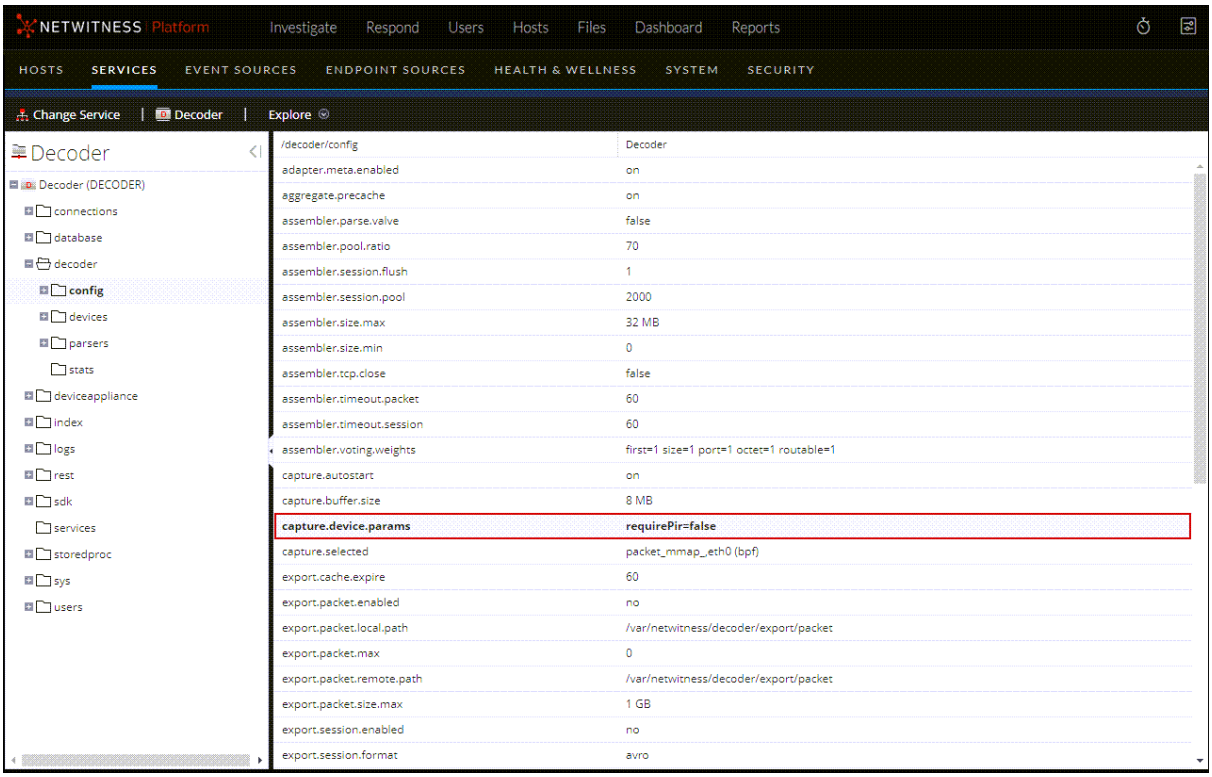
(Optional) Process Raw Syslog Data without Priority Field

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`



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

(Optional) Configure Decoder to Support OpenAppID

The NetWitness 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) 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
`client_*`	Detectors for client applications (for example, Facebook, MacAppStore).
`content_group*`	Detectors for a specific type of network attribute such as protocol, port, or service (for example, port_services which contains many services identified by their common ports).
`payload_*`	Detectors for the host based traffic (for example, Hulu, Reddit).
`service_*`	Detectors for service applications (for example, Minecraft, VNC).
`ssl_host_group_*`	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
`client_9P.lua`	loaded
`client_9p.lua`	not loaded (Linux only)
`client_9P`	not loaded
`client_9P*`	loaded
`client_`	loaded

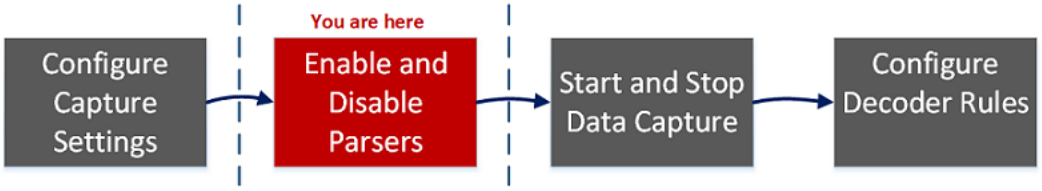
Detector Group Filename	Status
9P.*	
lient_ 9P.lua	not loaded
lient_ 9P.*	not loaded
*lient_ 9P.lua	loaded

5. Once the required detectors are enabled, reload the parser to load the detectors. The following example message will confirm that the detectors and 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 must 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.

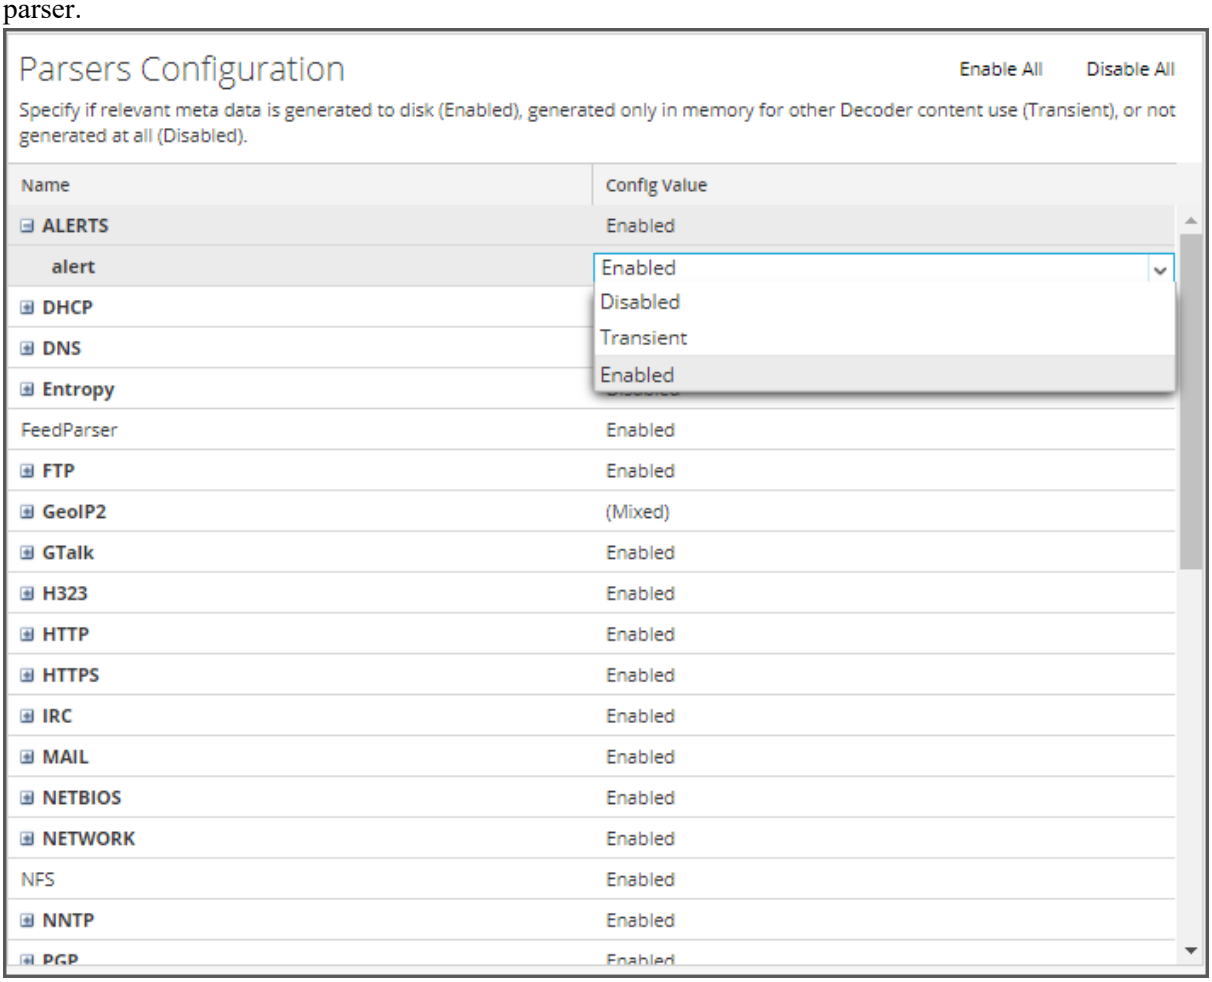
The Parsers Configuration panel provides options to select parsers to use on the Decoder. Within some parsers, you can also configure the metadata that the parser creates. These are the options in the Parsers Configuration panel.

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 *Data Privacy Management Guide*). • 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 NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

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.



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. The status for each parser is updated, based on your selections.

(Optional) Decoder Parser Utilization Insights

To make insights on parser usage more accessible and useful, NwDecoder and NwLogDecoder were updated to expose the `/decoder/parsers/definitions` folder and a number of stats within the parsers that give some detail on utilization. Previously, the folder, parser folders within, and curated parser stats were gatekept by `detailed.stats`.

The parser stats that were exposed were the following:

- Memory usage.
- Callbacks for metas, ports, and tokens.
- Callback counts for metas, ports, and tokens.

Configuration and Usage

In addition to being stored in the original statdb when configured to do so (default), the newly exposed stats are being stored in a separate database in the same logical volume (`decodersmall`) as the original stats database, `parsestatdb`. The content of parsestatdb is limited to the parser folders stored in `/decoder/parsers/definitions` at this time.

By default, `parsestatdb` is configured for 10 GB and the decodersmall volume was extended by 10 GB to accommodate the new directory.

For all intents and purposes, parsestatdb works the same as the original statdb in terms of functionality, just with some usage items specific to it + a new config item that controls how often to update the newly exposed stats (applicable to original statdb as well):

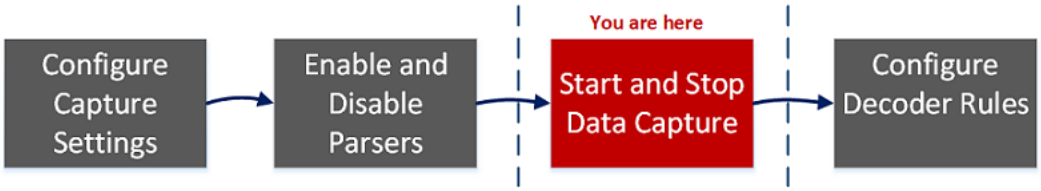
- `/decoder/parsers/config/parser.stat.dir`: Specifies where curated parser `statdb` files should be located. When you modify the size, the change takes effect immediately. However, when you add or remove directories, the changes take effect when the service restarts.
- `/decoder/parsers/config/parser.stat.compression`: Enable to compress parser stats as they are written to the database. Compression types available are none, gzip, bzip2, lzma and zstd. Change takes effect immediately.
- `/decoder/parsers/config/parser.stat.exclude`: A comma separated list of stat pathnames to be excluded from the parse stat db (meta and CPU stats not included at this time). The following wildcards are permitted: ? match any single character, * match zero or more characters to delimiter /, ** match zero or more characters including delimiter. Change takes effect immediately.
- `/decoder/parsers/config/parser.stat.interval`: Determines how often (in milliseconds) detailed parser statistic nodes are updated. Cannot be faster than sys stat updates. Change takes effect immediately.
- `/decoder/parsers statHist`: Retrieve historical stats from the stats db. Do not send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char, * to match zero or more characters, not including slash /, ** to match zero or more characters including slash /.
- `/decoder reset`: new parameter added to reset just parsestatdb that only applies to Decoders.

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**.

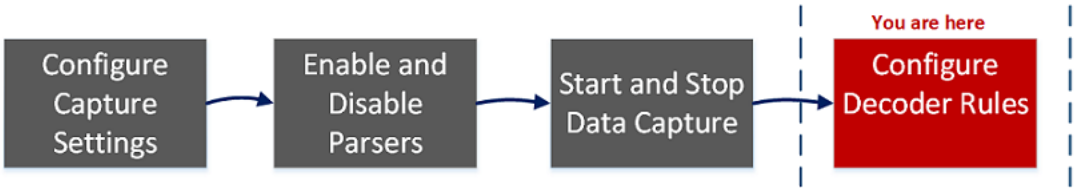
   Packet or log capture ceases, and the option to upload a file to the service is again available.

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 Core Services must follow.

By default, no rules are defined when you first install NetWitness. 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 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 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, see "`WHERE` Clauses" in the "Queries" section of the *Core Database Tuning Guide* .

Rule Name	Condition
ComplianceDevices	`device.group='PCI Devices' || device.group='HIPPA Devices'`
HighValueWindows	`device.group='Windows Compliance'`
MediumValueWindows	`device.type='winevent_nic' && msg.id='security_4624_security'`

Rule Name	Condition
LowValueWinLogs	`device.type='winevent_nic' && msg.id='security_4648_security'`
LowValueProxyLogs	`device.class='proxy' && msg.id='antivirus_license_expired'`
GeneralWindows	`device.type='winevent_nic'`

Invalid Rules

NetWitness 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 logs indicating the error.

Note: NetWitness does not support parsing of legacy syntax rules. 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 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	`<=`	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	`>`	compatible with left operand	Greater than operator. You can use values or meta keys on the right side of this operator.
any	`>=`	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	`contains`	text	Find values that contain the right operand. You can use meta keys or values on the right side of this operator.
text	`begins`	text	Find values that begin with the right operand. You can use meta keys or values on the right side of this operator.
text	`ends`	text	Find values that end with the right operand. You can use meta keys or values on the right side of this operator.
text	`length`	integer	Find strings of a certain length. You can use meta keys or values on the right side of this operator.
any	`count`	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	`ucount` and `unique`	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 `ucount` is six.
N/A	`exists`	any	Finds any values for the meta key. You can use meta keys or values on the right side of this operator.
N/A	`!exists`	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	`regex`	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: `tcp.port = 40000-u`
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: `tcp.port = l-40000`
- (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: `tcp.port = 25-443`
, (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: `tcp.port = 1-10,25,110,143-225,40000-u`
( )	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: `(alias.ip=192.168.1.1 && tcp.port=80) || (alias.ip=10.10.10.1 && tcp.port=443)`
~	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

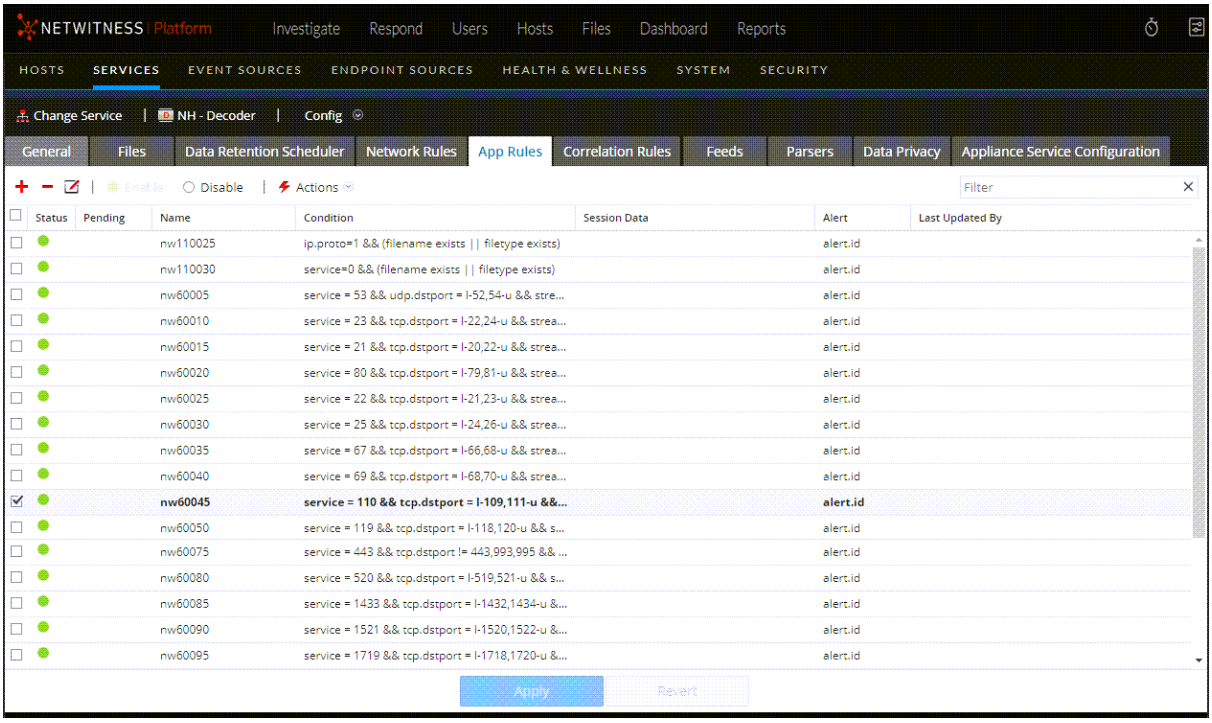
The Decoder and Log Decoder rules are editable in the Services Config view. While each type of rule (network, application, and correlation) has its own tab; the functions are similar for all types of rules. You can:

- 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.



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.

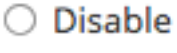
To remove a rule:

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

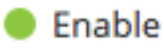
1. From any Rules tab, select the rule to edit.
2. Click  or double-click the rule row.
   The Rule Editor dialog for that type of rule is displayed.

To disable a rule:

1. From any Rules tab, select the rules to disable.
2. Click 

   .
   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 

   .
   The status changes to enabled in the rules list, but the rule is still disabled on the service.

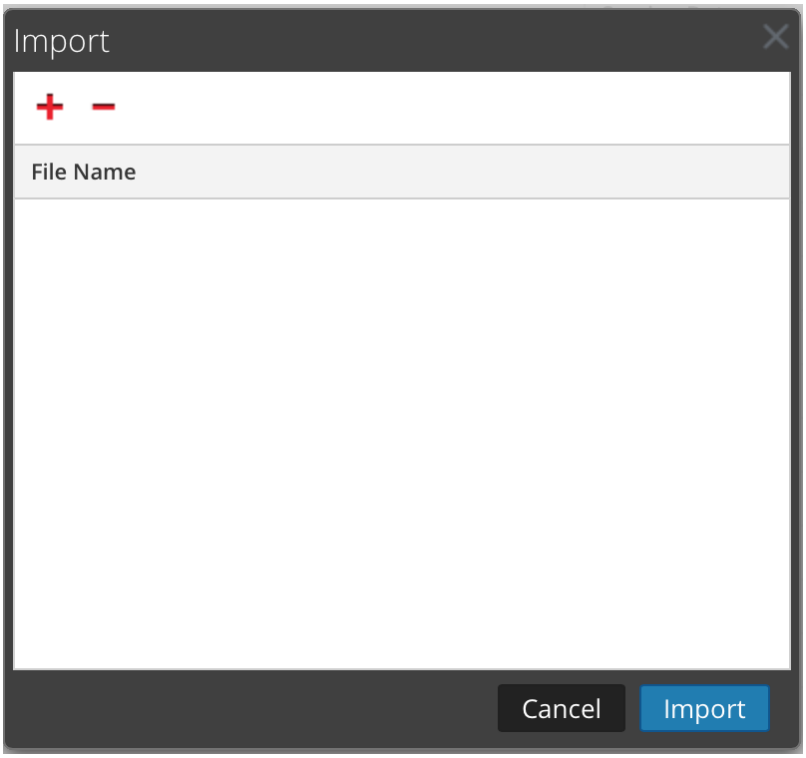
Import Rules from a File and Export Rules

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.

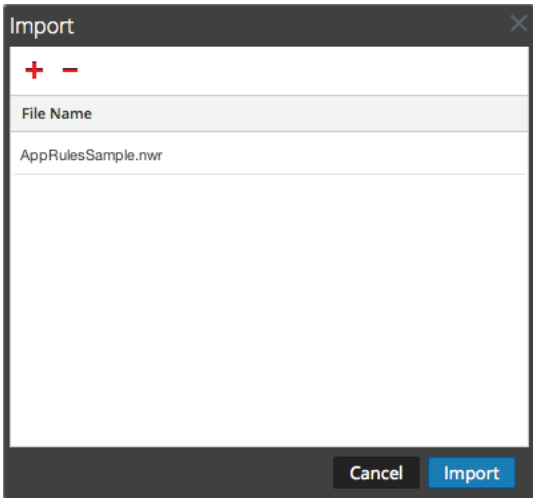
When you attempt to import a group of rules, NetWitness Administration checks the type of rules imported. If you are successful, a message displays the number of rules imported. If the rule type differs from the active tab type, the rules are not imported. You must re-import the rules under the correct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

To import rules to a service:

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

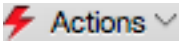


2. Click + .
   A view of the directory structure is displayed.
3. Choose one or more NetWitness rules (.nwr) files to import, and click **Open**.
   The file is added to the list in the Import dialog.



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 

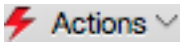
 > Export > Selection. (**Export > All** exports all rules in the rules list even if you have a subset selected for export.)
   - Right-click the selected rules and select **Export Selection**.

     A prompt for the filename is displayed.
3. Enter the filename and click **Export**.
   The .**nwr** file is downloaded.

Push Rules to Other Services

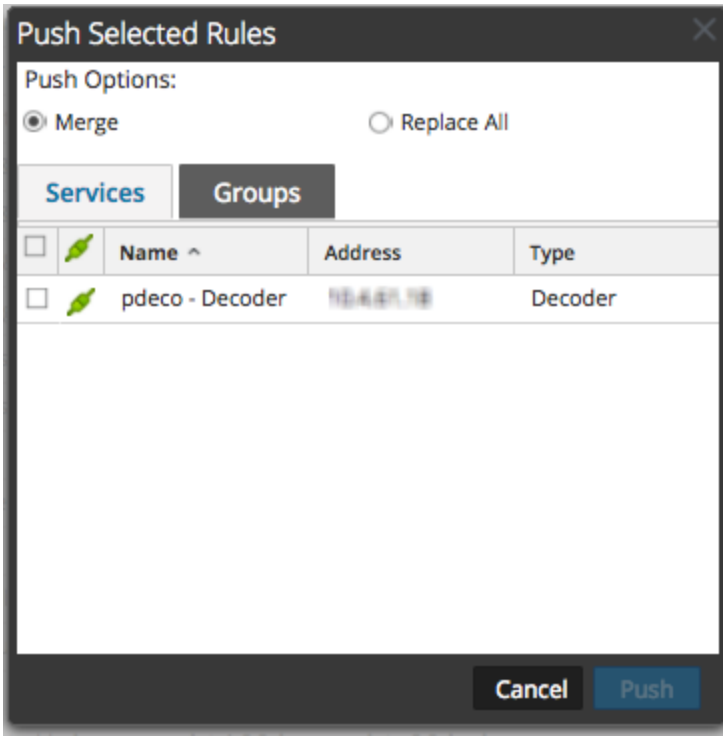
You can apply (push) rules or selected rules to other services (Decoders or Log Decoders) or service groups. When you push all rules to other services, all rules on the target services are removed and replaced with all of the rules on the source service.

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 

 > Push > Selection.

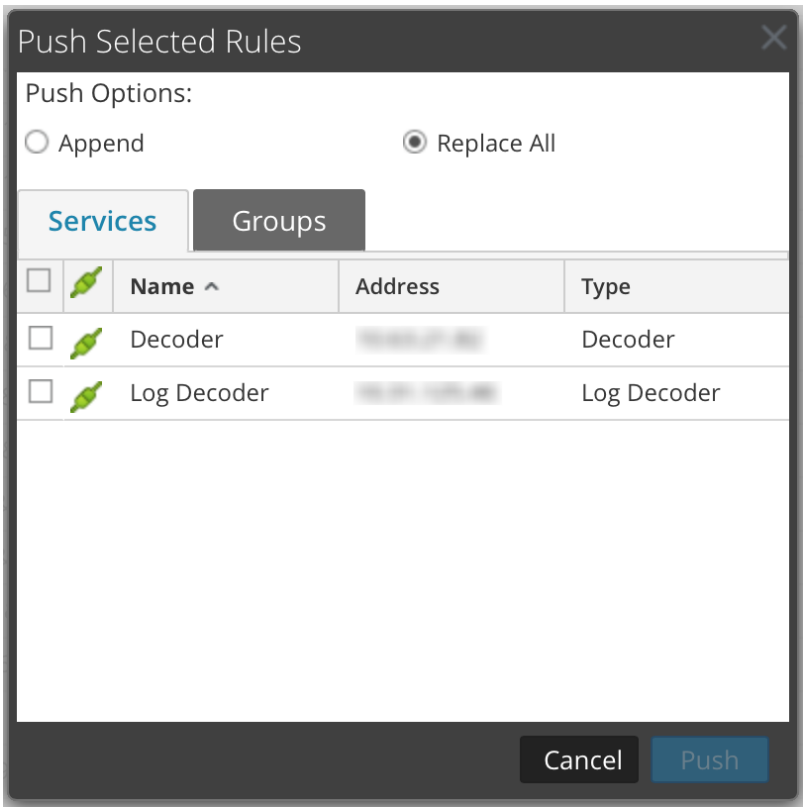
- Right-click the selected rules and select **Push Selected Rules**.
  The Push Selected Rules dialog is displayed.



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.



2. On the **Services** tab, select the target services to receive the pushed rules, or select the groups of services from the **Groups** tab.
3. 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:

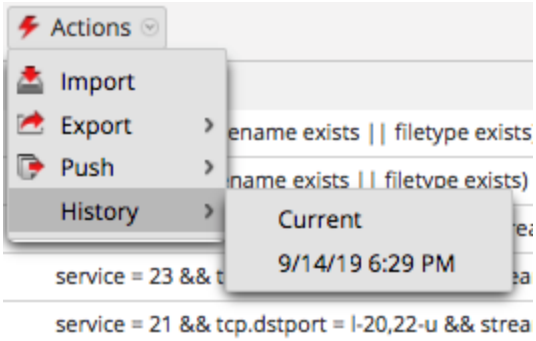
- Drag and drop the rules in the appropriate location in the rules list.
- Right-click a rule to display the context menu, and use the **Cut** and **Paste** options.

Restore a Rule Snapshot from History

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

To restore a rules snapshot from history:

1. 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:

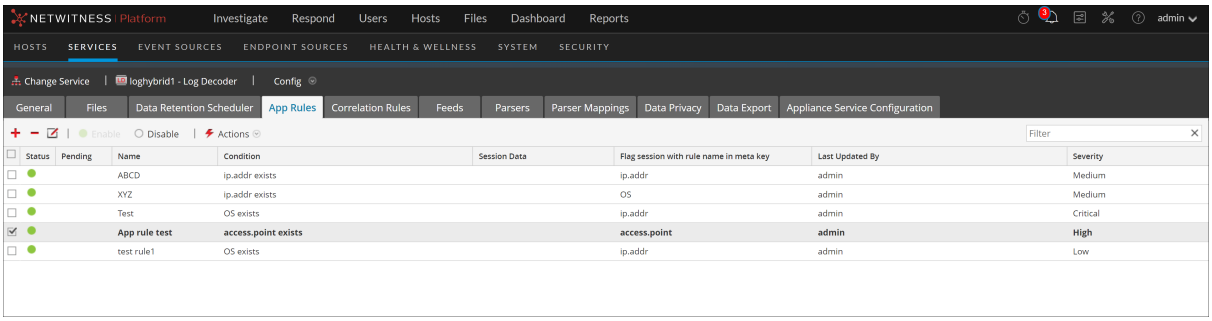
- Rule Name: Truncate SMB
- Condition: `service=139`
- Rule Action: Truncate All

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

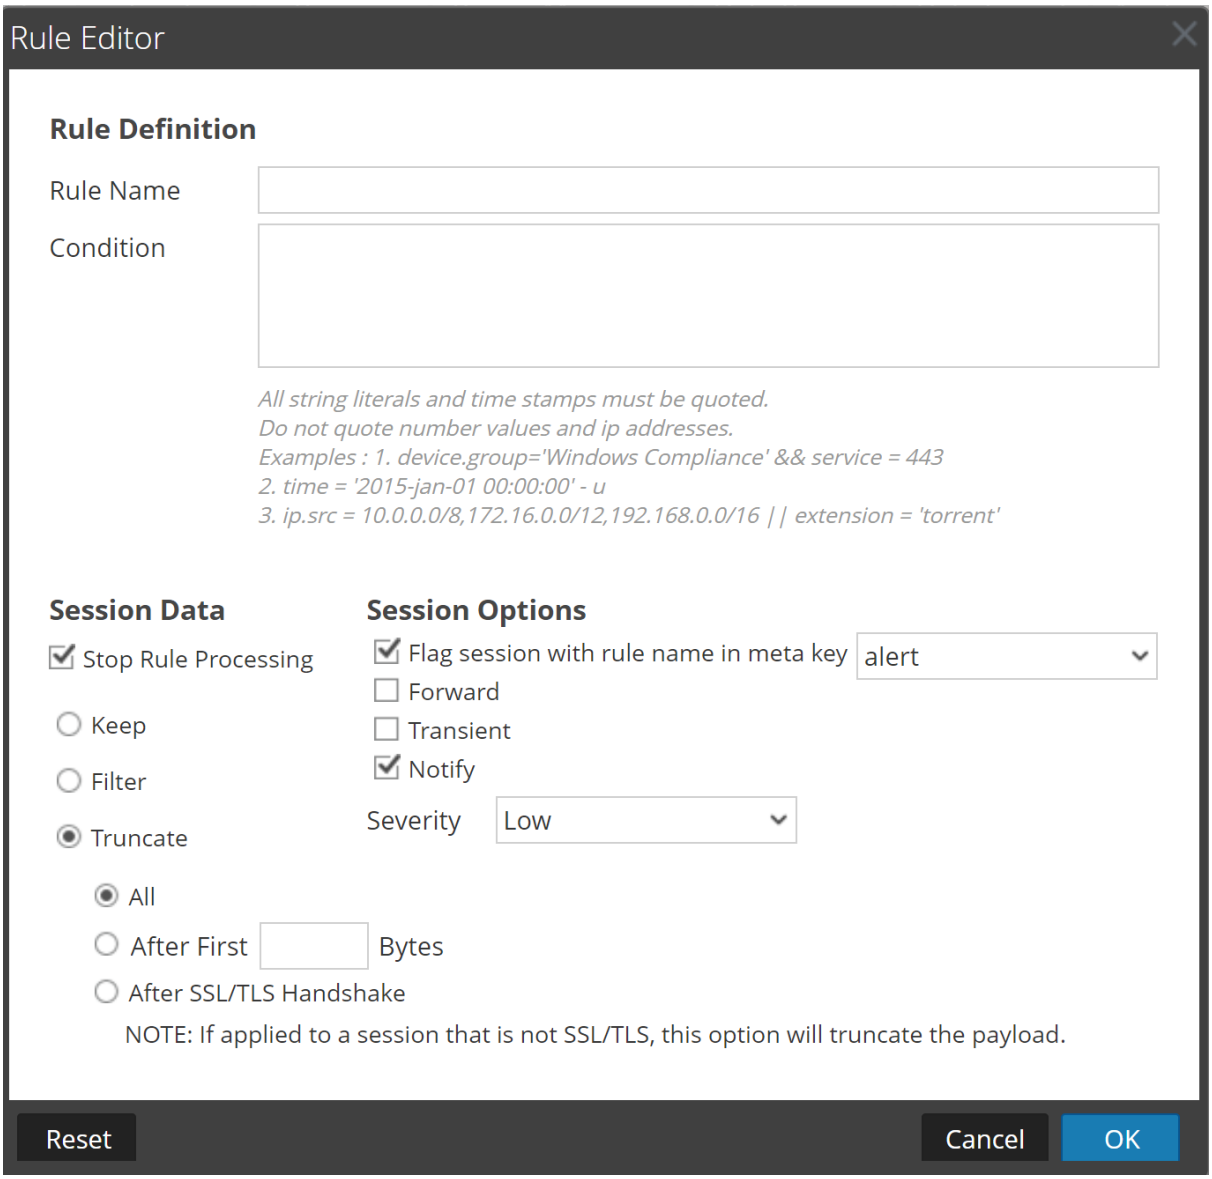
- Rule Name: Retain Email Tom Jones
- Condition: `email='Tom.Jones@TheShop.com'`
- Rule Action: Keep

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.



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.



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 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:

- Enable the **Flag session with rule name in meta key** and select the metadata for the alert from the drop-down list, which generates a custom alert when metadata matches the rule. This is mandatory.
- **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.
- Enable the **Notify** option to choose the **Severity** level from the drop-down menu for the application rule created. The available options are listed below:
  - Low
  - Medium
  - High
  - Critical

  Note: Severity is selected by default as **Low**.

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 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 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

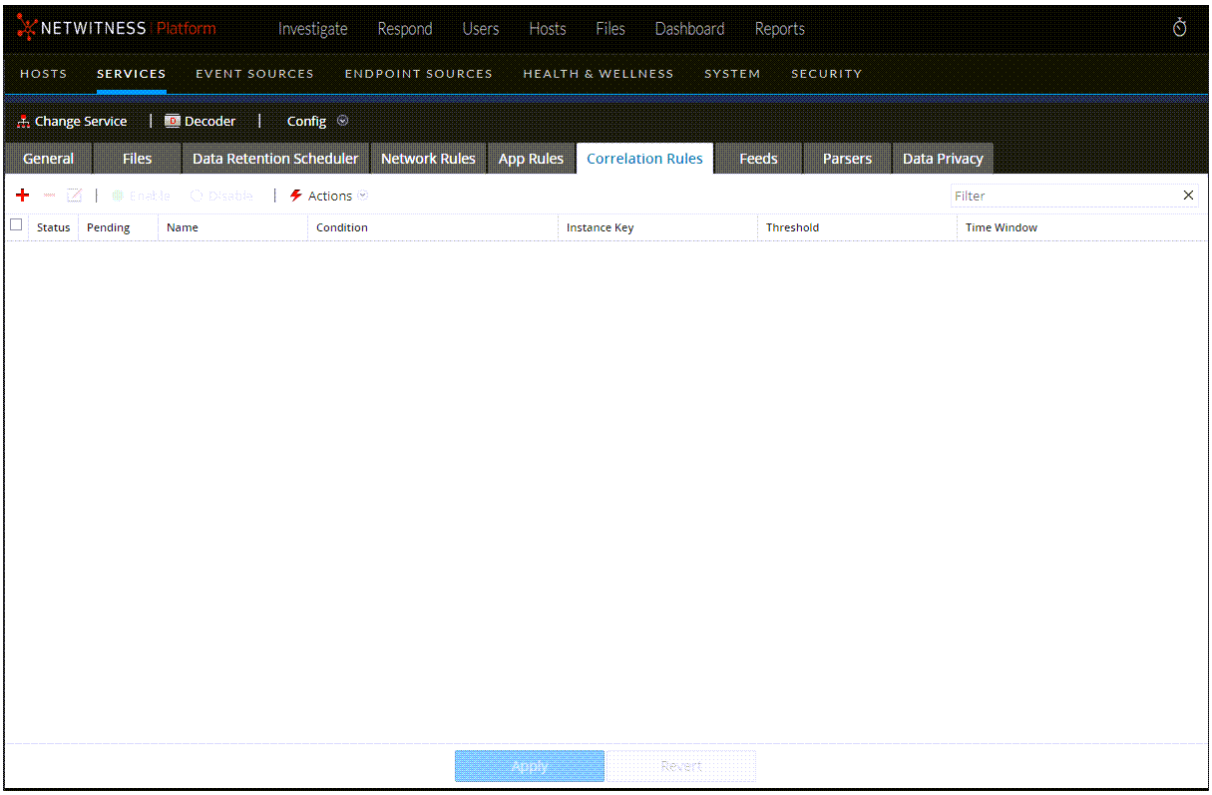
Both sample rules have the same instance key: `ip.src` and `ip.dst`. Because we are looking for unique combinations of `ip.src` and `ip.dst` that match the correlation condition, **`ip.src`** and **`ip.dst`** are **primary keys**.
Threshold can include an **associated key** that identifies the meta type 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.

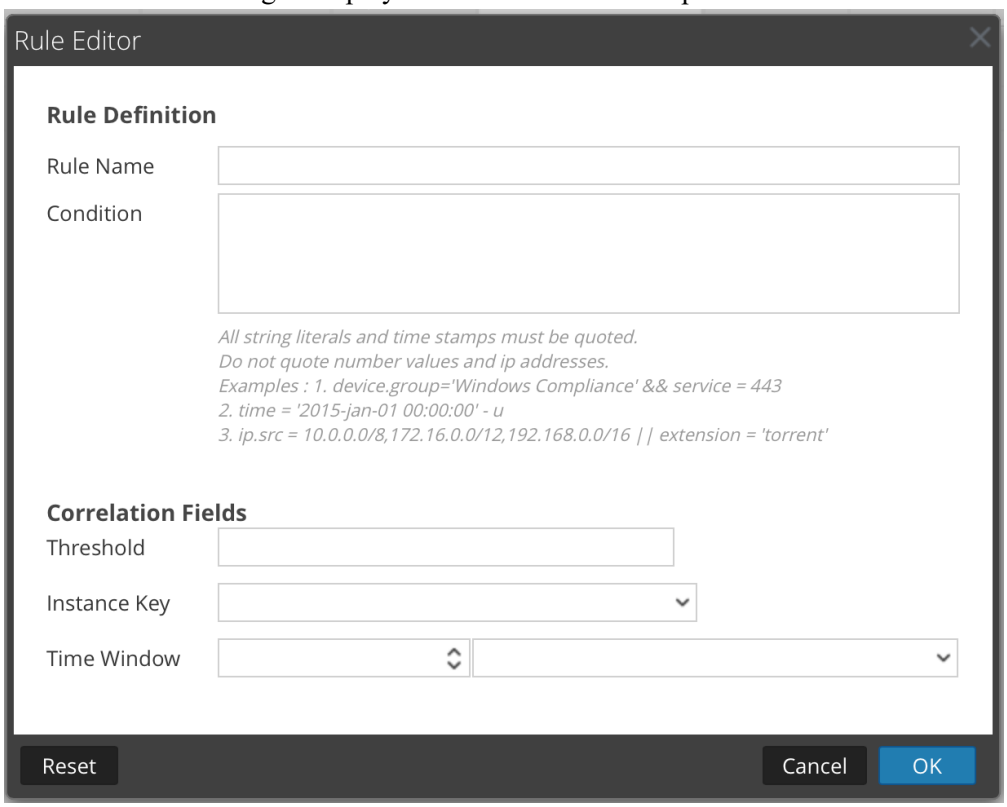
2. Select the **Correlation Rules** tab.



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.



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. 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 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 55553, the condition should be written as follows:
`port.src=553 || port.dst=553 || port.src=55553 || port.dst=55553`
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 supports for use in network rule conditions.

Meta Key	Description
`eth.addr`	Ethernet source or destination address. Commonly known as the MAC address.
`eth.dst`	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).
`eth.src`	Same as Ethernet destination except that it focuses on the source address.
`eth.type`	Ethernet frame type.
`hdlc.type`	Frame type of the HDLC frame.
`alias.ip`	IPv4 source or destination address in standard form. IP addresses can be entered in CIDR notation for subnets.
`ip.dst`	Destination IPv4 address in standard form. IP addresses can be entered in CIDR notation for subnets.
`ip.proto`	IPv4 protocol field.
`ip.src`	Source IPv4 address in standard form. IP addresses can be entered in CIDR notation for subnets.
`alias.ipv6`	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
ipv6.dst	Destination IPv6 address in hex format.
ip.proto	IPv6 protocol field. This maps to the Next Header field in the IPv6 header and uses the same values as the IPv4 protocol field.
ipv6.src	Source IPv6 address in hex format.
port.dst	Destination TCP port.
tcp.port	TCP source or destination port.
port.src	Source TCP port.
port.dst	Destination UDP port.
udp.port	UDP source or destination port.
port.src	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:

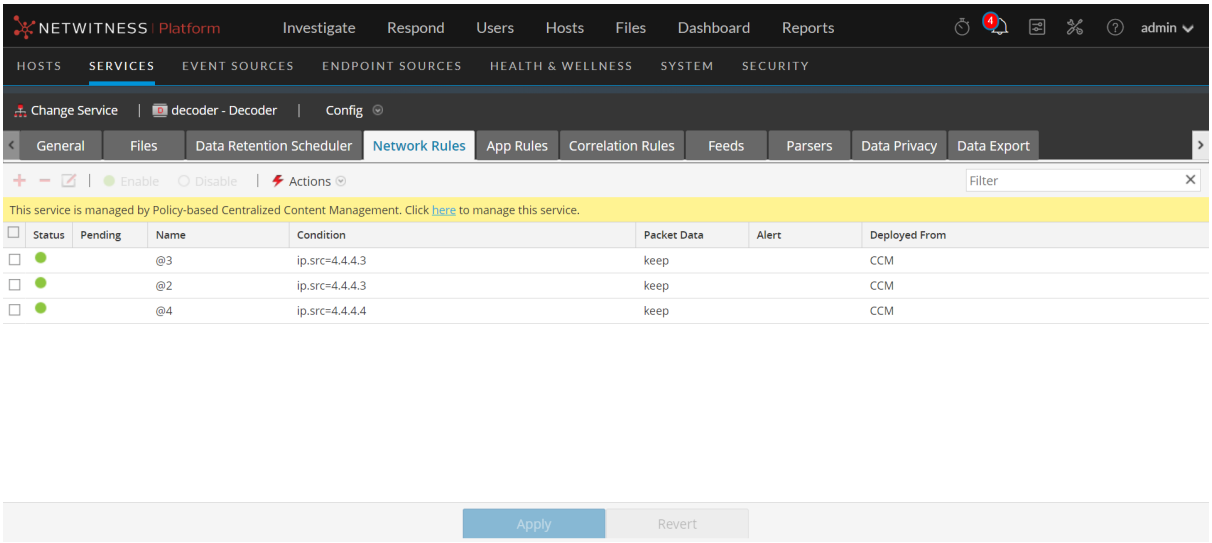
- Rule Name: Subnet Filter
- Condition: `alias.ip=192.168.2.0/24`
- Rule Action: Filter

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*.

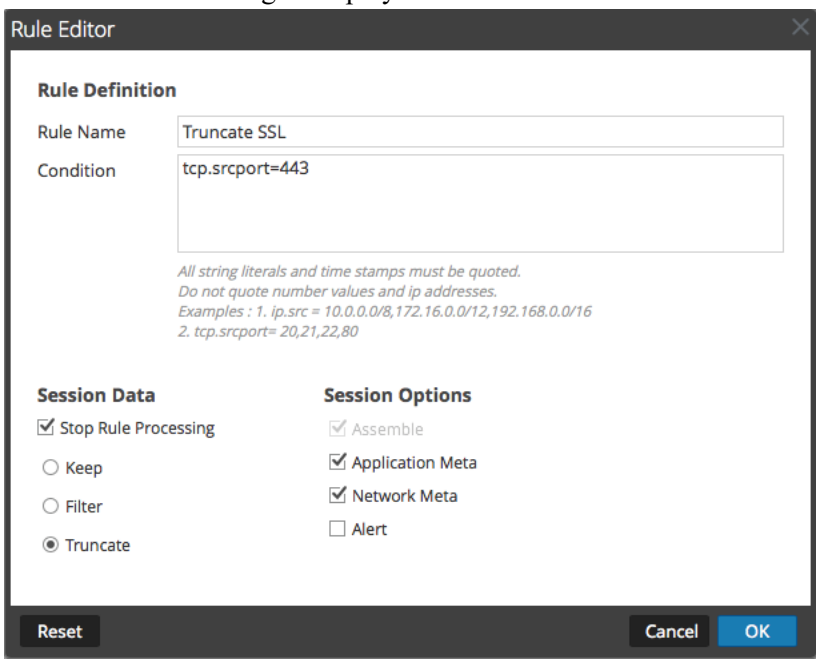
To add or edit a network rule:

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.



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 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 supports for use in network rule conditions.

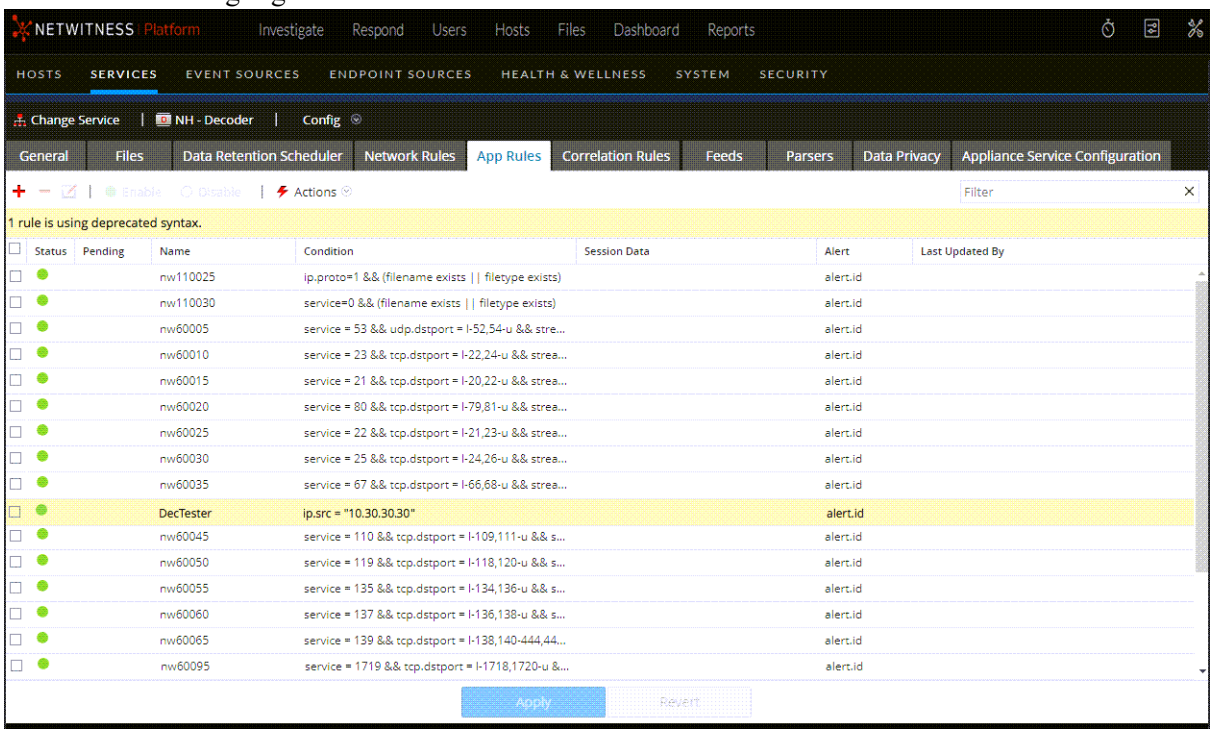
6. If you want rule evaluation to end with this rule, select the **Stop Rule Processing** checkbox.
7. In the **Session Data** section, choose one of the following actions to apply when a matching packet is found:
   - **Keep**: The packet payload and associated metadata are saved when they match the 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 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

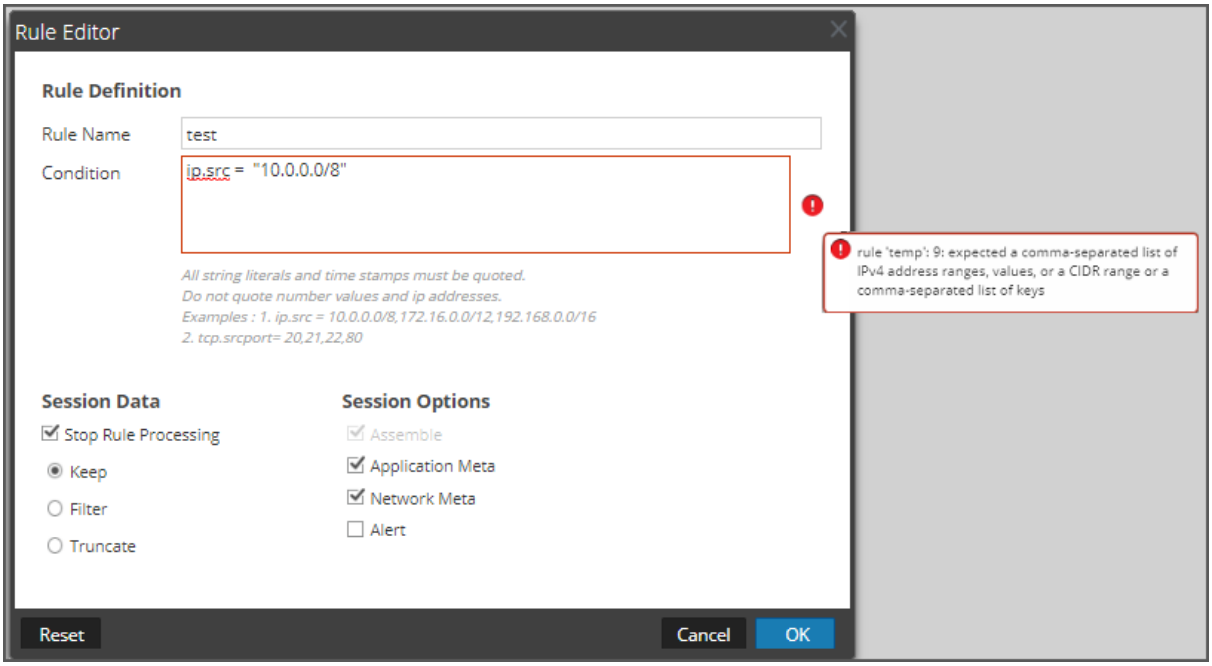
After an update to NetWitness latest version, the user interface highlights any rules with invalid syntax. The Rule Editor provides additional tooltips. After you fix the rules, the highlights disappear. Configure Decoder Rules provides guidelines that all queries and rule conditions in NetWitness must 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.



4. Select an invalid rule and click  .
   The Rules Editor shows additional information for the invalid rule.



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 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 NetWitness 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%20n
ame%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 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, 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 Live Services.

Upload and Delete Custom Parsers

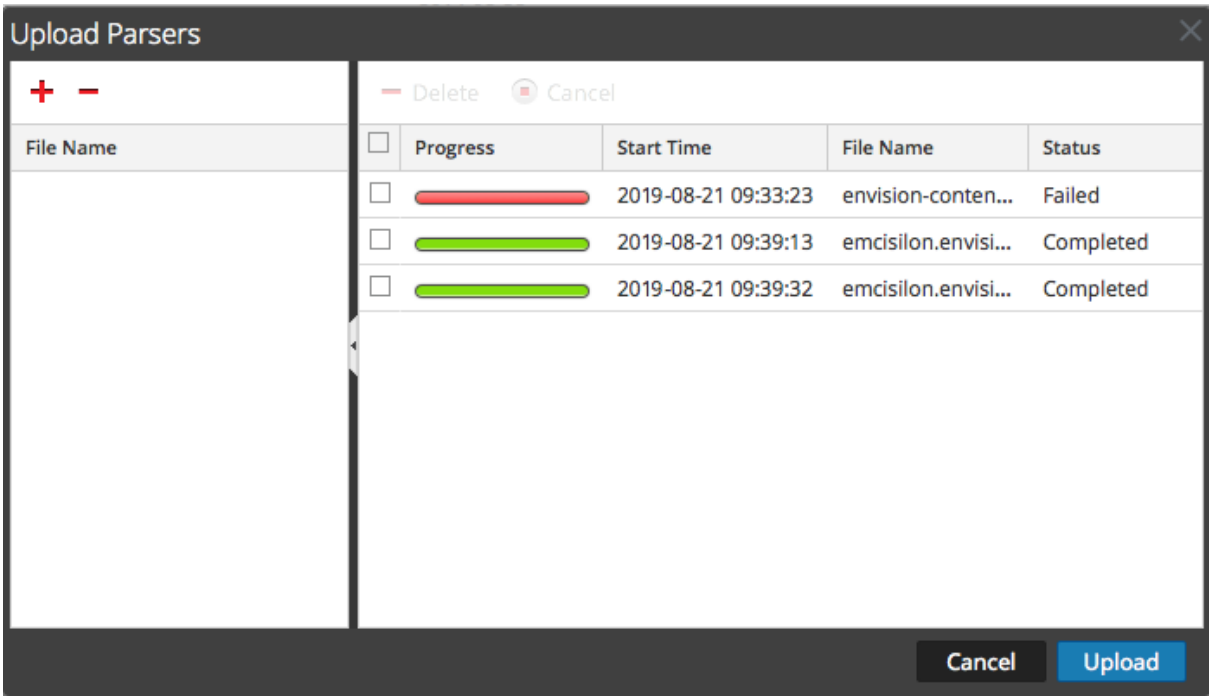
NetWitness has the ability to upload parsers from your local system and delete these parsers.

Upload Parsers to a Decoder or Log Decoder

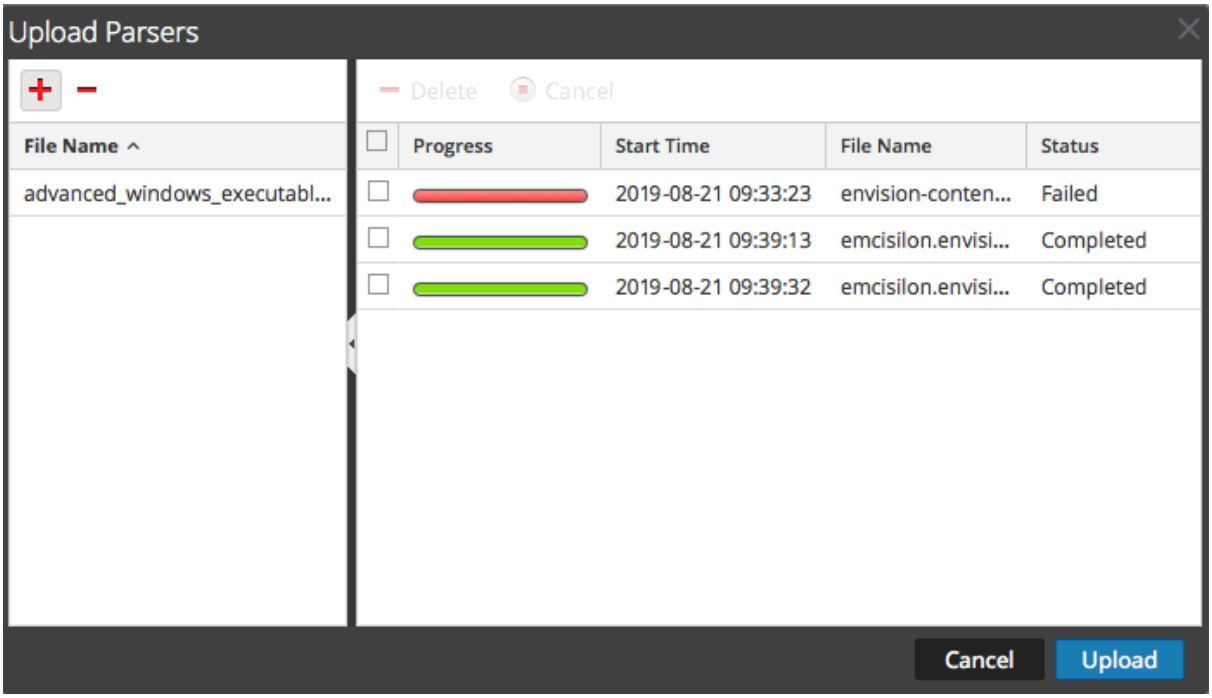
The Upload option in the Service Config view > Parsers tab displays the Upload Parsers dialog, in which you can manage the uploading of parsers to a Decoder or Log Decoder. In the File list, you prepare a list of parsers for uploading. You can add files from a directory structure, and delete files from 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.

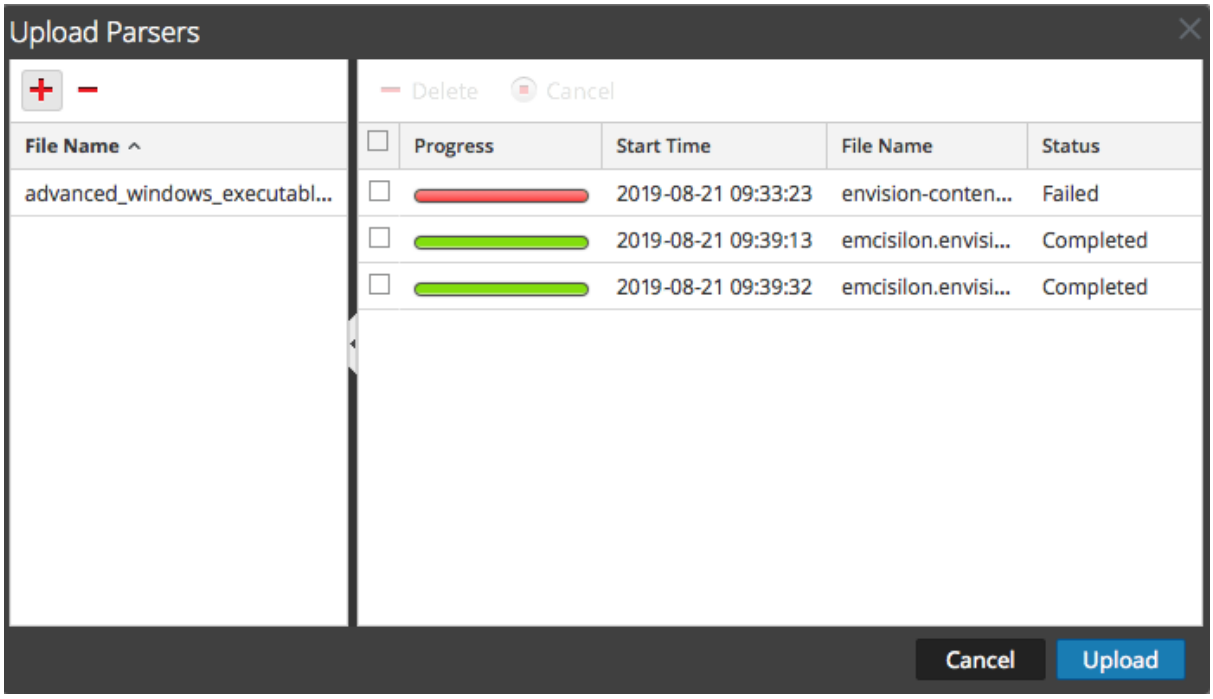
3. Click Upload.
   The Upload Parsers dialog is displayed.



4. Click .
   A file selection dialog is displayed.
5. 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 .

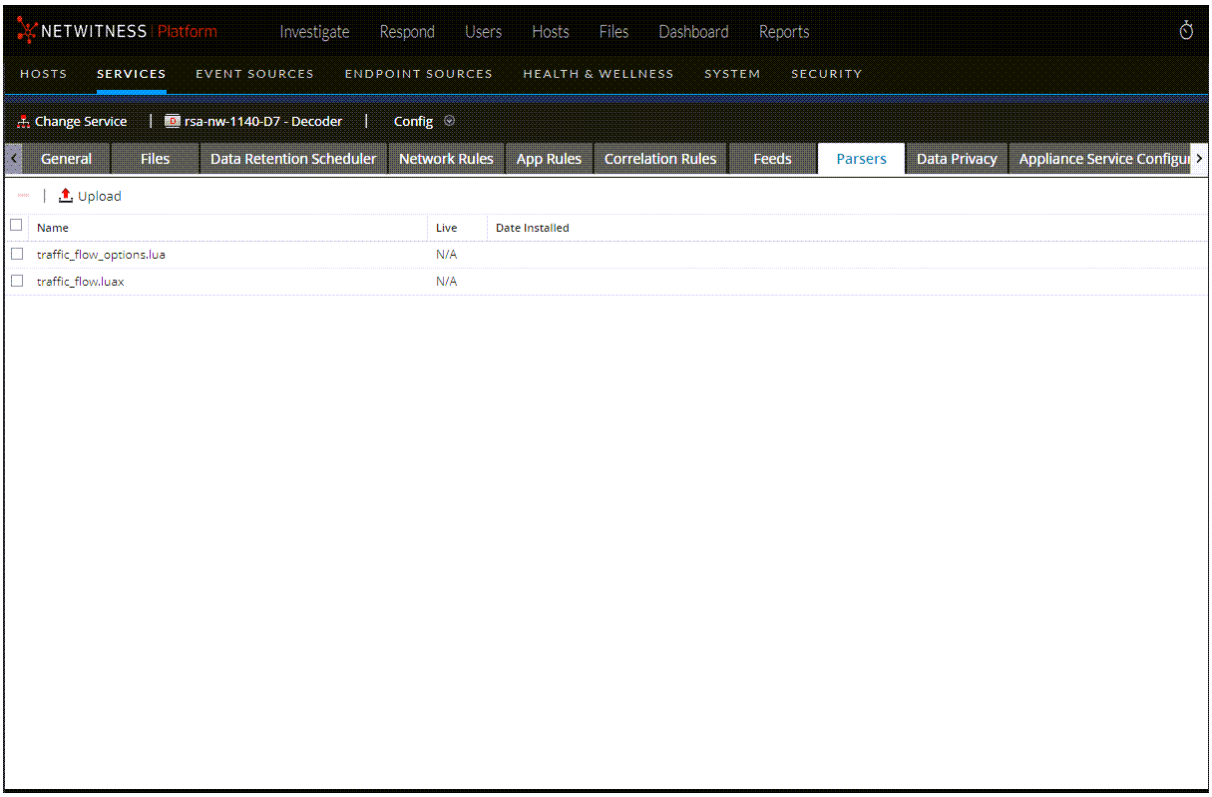
Delete Deployed Parsers

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

In NetWitness Platform, 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 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`.

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

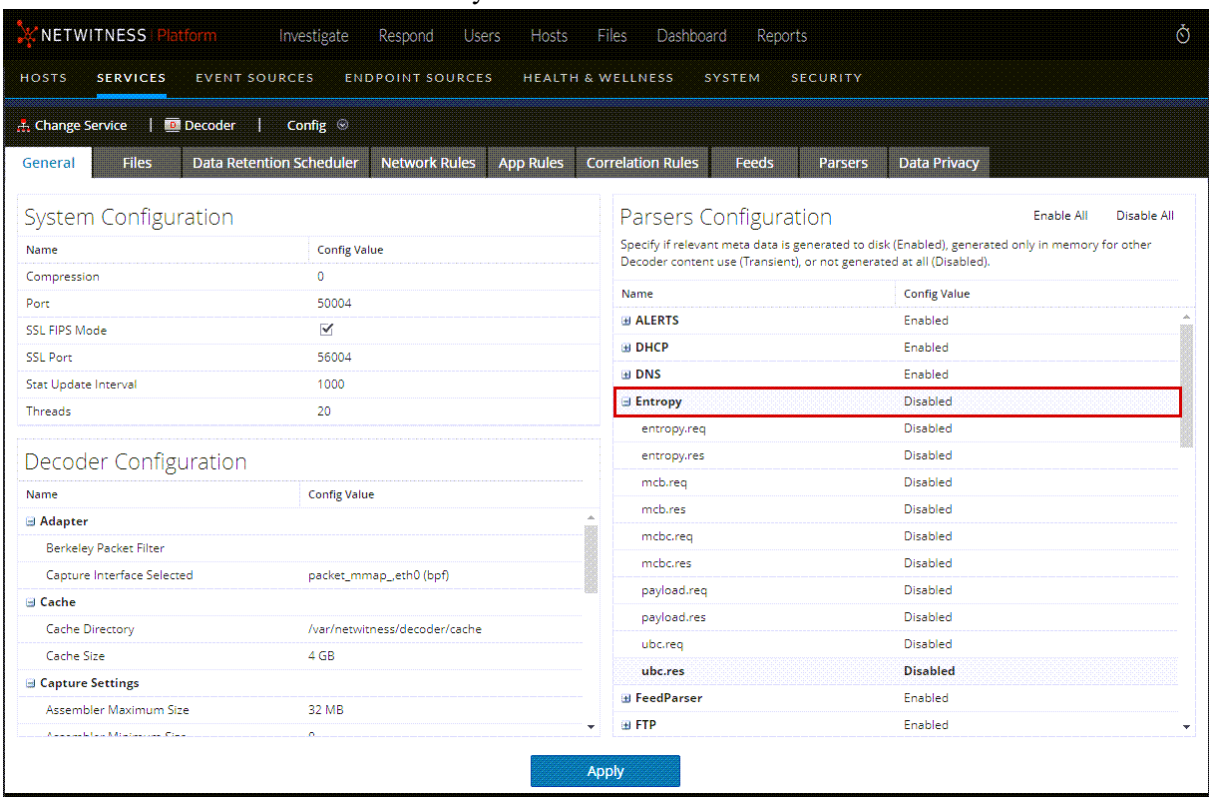
These are the new meta keys generated by the Entropy parser by default:

- `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 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.

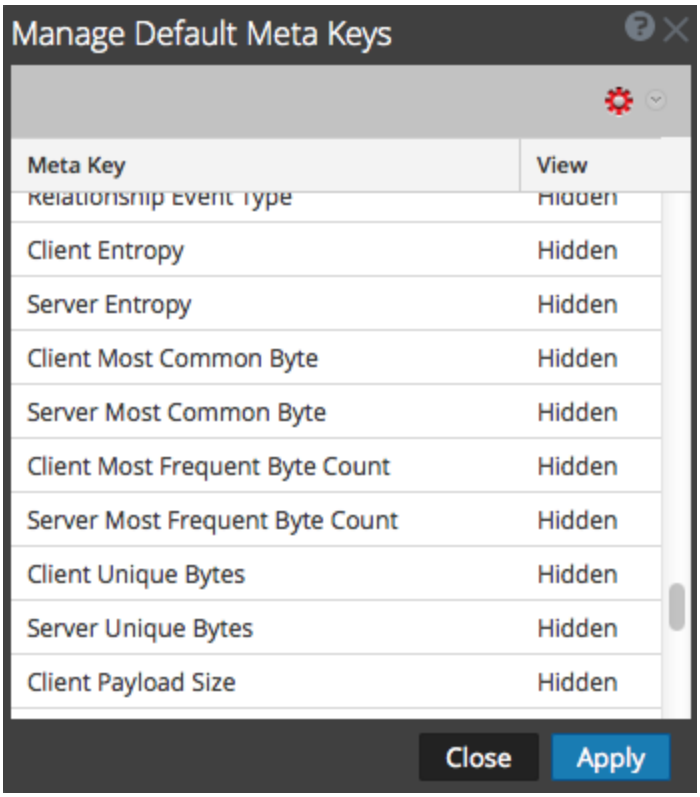
To enable and configure the Entropy parser on a Decoder:

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.



3. Click **Apply**.
   The Entropy parser is enabled and begins creating the new meta keys as configured in the Concentrator custom index file.
4. 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.
5. 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 inital 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: mstshash=" />
          </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>
```

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
`parsers`		The root node in each definition file.
	`xmins:xsi`	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
	`xsi:noNamespaceSchemaLocation`	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: parsers.xsd
`parser`		The node that defines a single parser definition. This node must be directly under the `parsers` node. There can be more than one per file.
	`name`	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].
	`desc`	Provides a friendly description of what the parser does.
	`service`	The unique number assigned to the session when identified.
`declaration`		Delineates the definition. Each of these definitions can have an associated `match` entry.

Node Name	Attribute Name	Description
token		Specifies a definition for identifying a token somewhere in the session protocol. This defines a match callback when the specified tokens are encountered in a session payload. The read 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 (linestart or linestop).
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 apptype (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 apptype. (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 `match` 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 `regex` 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 `session` (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 `constant`.
`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
`find`		Searches the stream payload starting at the read position for a provided string value. If the value is found, the offset from the read position is returned. Any child elements will then execute. If not found, any child elements will not execute.
	`name`	A `number` variable to receive the offset from the `read` position where the match begins.
	`value`	A string 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.
`install-decoder`		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.
	`type`	The type of decoder to install. Valid options are: gzip, deflate, chunked, chunked-gzip, chunked-deflate.
	`offset`	Offset from the current read position to begin decoding.
	`length`	The maximum payload length to decode.
`isdecoding`		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 `read` position forward in the current stream by a specified number of bytes. If there is sufficient data in the stream, the `read` position is updated and any child elements will then execute. If not found, the `read` position remains unchanged and any child elements will not execute.
	`value`	The number of bytes to move the `read` position.
	`direction` (optional)	The direction to move the current read position. Can be `forward` (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 `read` position into a variable. If there is sufficient data in the stream, the `read` position is updated, the data read assigned, and any child elements will then execute. If not found, the `read` position remains unchanged and any child elements will not execute.
	`name`	The name of a `string` or `number` variable to receive stream data. If a `number` 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 `big` (default) or `little`. The attribute is invalid when reading into a `string` 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 `number` variable to receive the offset from the `read` 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 `string` 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 `string` 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 `number` variable to receive the zero-based position, where the provided value string was found in the **in** string.
	`value`	A string to find.
	`in`	A string to search.
	`length` (optional)	A limit to the length of the **in** string to be searched. If a limit is not provided, all of **in** will be searched.
`length`		Assigns the length of a string to a `number` variable.
	`name`	A `number` 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 number variable to receive the zero-based position, where the provided regular expression matched in the in string.
	`value`	A regular expression to be searched for.

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

In NetWitness latest version, 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 latest version. 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 NetWitness 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. 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 parsers.options
`true`	Decompress all bodies	`HTTP="decompress=true"`
`false`	Do not decompress. This is the default.	`HTTP="decompress=false"`
`1`	Decompress `application/*` content	`HTTP="decompress=1"`
`2`	Decompress `audio/*` content	`HTTP="decompress=2"`
`4`	Decompress `font/*` content	`HTTP="decompress=4"`
`8`	Decompress `image/*` content	`HTTP="decompress=8"`
`16`	Decompress `message/*` content	`HTTP="decompress=16"`
`32`	Decompress `model/*` content	`HTTP="decompress=32"`
`64`	Decompress `text/*` content	`HTTP="decompress=64"`
`128`	Decompress `video/*` content	`HTTP="decompress=128"`

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. Expland **decoder** > **parser**s 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 and data in the HTTP/2 stream to gain visibility into HTTP/2 sessions. The HTTP/2 parser supports demultiplexing interleaved streams and extracts the application payload for detections in other parsers looking at the tokens in the payload.

To turn on header parsing and decompression for HTTP/2 sessions:

1. Go to  (Admin) > **Services** and select a Decoder, and in the actions menu (), select **View** > **Explore**.
2. Expland **decoder** > **parser**s and select **config**.
3. In **parsers.options**, append any option shared below:
   a. For header only parsing.
   Append `HTTP2="headers=true"`
   b. For header parsing and header extraction to other parsers (example: HTTP_lua).
   Append `HTTP2="headers=true extract=headers"`
   c. For header parsing, header extraction and data decompression (example: gzip and brotli encodings).
   Append `HTTP2="headers=true extract=all"`

   Note: On enabling data decompression with `extract=all` the capture rate reduces by 10%.
4. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Note: On content reconstruction HTTP/2 compressed streams are automatically detected and decompressed.

To avoid duplicate meta from HTTP/2 parser when Lua parser HTTP_lua is enabled:

1. Disable the HTTP/2 parser.

   To disable parser refer to the topic: Enable and Disable Parsers and Log Parsers
2. Go to  (Admin) > **Services** and select a Decoder, and in the actions menu (), select **View** > **Explore**.
3. Expand **decoder** > **parsers** and select **config**.
4. Make sure to append `HTTP2="headers=true extract=headers"` or `HTTP2="headers=true extract=all"` in **parsers.options**.

   Note: On enabling data decompression with `extract=all` the capture rate reduces by 10%.

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

This would make HTTP/2 force enabled but with meta generation disabled, thus avoiding duplicate meta generation.

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 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: NetWitness 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 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 railed 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.netwitness.com/t5/netwitness-platform-unified-data/introduction-to-the-rsa-netwitness-platform-unified-data-model/ta-p/565620. 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 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 `ruletype` 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 (`tcp`, `udp`, `icmp`, `ip`).
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
`nocase`	Case-insensitive pattern matching is used.
`offset`	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.
`depth`	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.
`distance`	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.
`within`	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 conjuction with AppHTTP native parser:

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

Option	Description
`http_method`	The token must appear in the HTTP method field. The detection engine does not currently restrict the method tokens, such as `GET` or `POST`, to the HTTP method field, so these matches will work even if `AppHTTP` is turned off.
`http_uri`	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
`to_client`	Limits the rule to only matching on a stream that a Decoder has defined as Server.
`from_server`	Synonym for `to_client`.
`from_client`	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 `pcre` 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 `pcre` 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. You can configure the parser by editing the **search.xml** 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. Currently, NetWitness supports only keyword searches (Search a stream for a specific set of words).

Search Methods

The Search parser uses case-sensitive keyword search method (search a stream for one or more exact, case-sensitive matches from a list of strings).

Syntax

For sample syntax, refer to the following image:

```
<search_config>
    <Maxrecon>SOME_INTEGER_GREATER_THAN_0</Maxrecon>
    <Maxsearch>0</Maxsearch>
    <search>
        <name>SOME_ALPHANUMERIC_NAME</name>
        <Services>LIST_OF_SERVICE_INTEGERS_SEPARATED_BY_SEMICOLONS</Services>
        <Keywords>LIST_OF_KEYWORDS_SEPARATED_BY_SEMICOLONS</Keywords>
        <MatchLimit>1</MatchLimit>
    </search>
    <search>
        <name>SOME_ALPHANUMERIC_NAME</name>
        <Services>LIST_OF_SERVICE_INTEGERS_SEPARATED_BY_SEMICOLONS</Services>
        <Keywords>LIST_OF_KEYWORDS_SEPARATED_BY_SEMICOLONS</Keywords>
        <MatchLimit>1</MatchLimit>
    </search>
    …
</search_config>
```

Note: The search name parameter should be alpha numeric with no spaces.

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. In NetWitness Platform, a 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 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 to obtain feeds from outside sources. "Live Content in NetWitness" in the *Live Services Management Guide* provides an overview of the Live content management tool.

Within the NetWitness 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 or manually. Feeds can be added, removed, and updated while a Decoder is running without affecting capture.

Custom Feed Definition File Structure

The NetWitness 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 are in the form `<filename>.feed`. To create a feed, NetWitness 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 can fetch the most current version of the file for each recurrence. After a NetWitness feed is created, you can download the feed to your local file system, edit the feed files, and then edit the NetWitness 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 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"
   valuetype="Text"
   apptype="0"
   truncdomain="true"/>

   <LanguageKeys>
            <LanguageKey name="threat.source" valuetype="Text" />
            <LanguageKey name="threat.category" valuetype="Text" />
            <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

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:

```
clc1,(bob;tom;sam) clc12,susan clc123,doris
```

This definition means that when we see the `lc.id` 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 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 Parameter	Feed Definition File Equivalent
(Define Feed Tab) Feed Type	Select: **Default** - to define a feed based on a `.csv` formatted feed data file. **STIX** - to define a feed based on STIX formatted `.xml` file.
(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, so NetWitness downloads a file at regular intervals and pushes it to the downstream devices.

NetWitness Parameter	Feed Definition File Equivalent
(Define Feed tab) **Name**	The custom feed name in the feed data file. It corresponds to the `flatfeedfile name` 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 `flatfeedfile path` attribute in the feed definition file. For example, `dynamic_dns.csv`.
(Advanced Options tab) **XML Feed File**	The name of the feed definition file. For example, `dynamic_dns.xml`.
(Advanced Options tab) **Separator**	The separator character used to separate attributes in the feed data file. It corresponds to the `flatfeedfile separator` 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 `flatfeedfile comment` 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 `MetaCallback apptype` attribute in the feed definition file. A value of 0 indicates no filtering by service type.
(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 `MetaCallback truncdomain` attribute. If the value is `www.example.com`, it is truncated to `example.com`. A value of **False** selects no truncation, and **True** selects truncation.

NetWitness Parameter	Feed Definition File Equivalent
(Define Columns tab, Define Index) **Callback Keys**	For a Non IP index, the available meta keys to match on instead of ip.src/ip.dst (the defaults for IP index type) are selectable from the drop-down list. The Callback Key corresponds to the `MetaCallback name` attribute, and the index column of the csv file must contain data that can match the chosen meta key. For example, if the `user` meta key is chosen, the index column of the `.csv` 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 `LanguageKey`, as defined in the feed definition file, for which meta is created from this row of the feed data file. It corresponds to the `Field key` attribute in the feed definition file. A key applies only to a field whose type is set to `value`. In the feed definition file, there is a list of `LanguageKeys` from `index.xml`, or a summary name if Source Name and Destination Name are used. For example, `reputation` is a summary name for `reputation.src` and `reputation.dst`. This value is referenced by the Field key 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.

`.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>
```

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

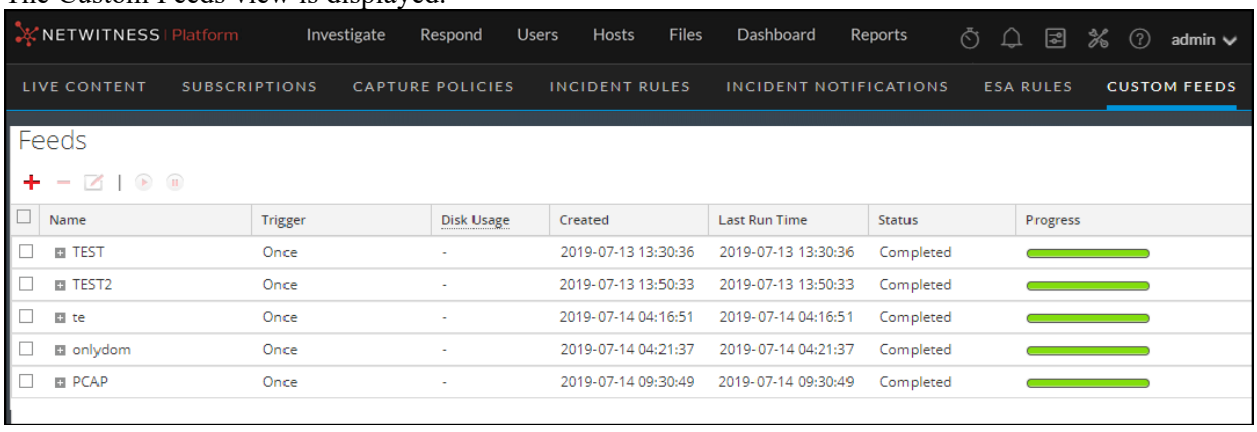
The feed data file and optionally the feed definition file (`.xml`) must be available on the local file system for an on-demand custom feed. For a recurring custom feed, the files must be available at a URL that is accessible to the NetWitness server.

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:

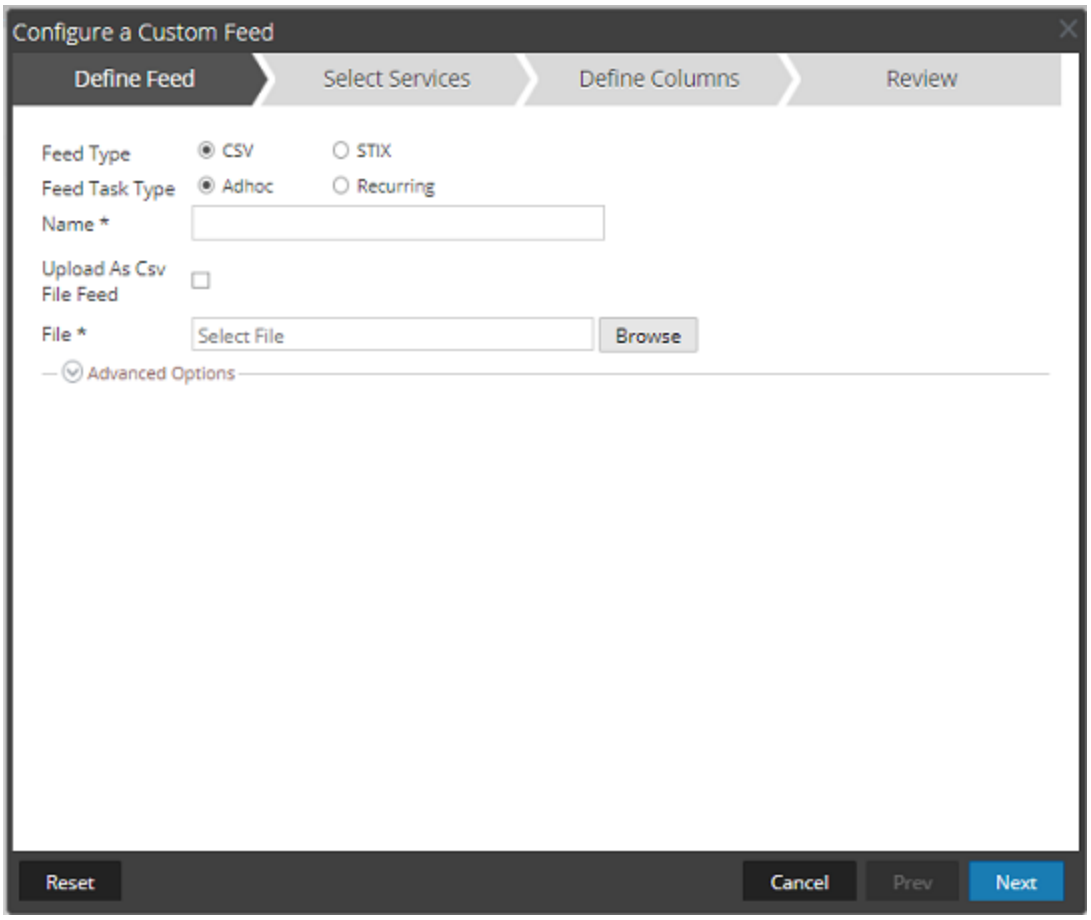
1. Go to **Configure > Custom Feeds.**

   The Custom Feeds view is displayed.



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

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



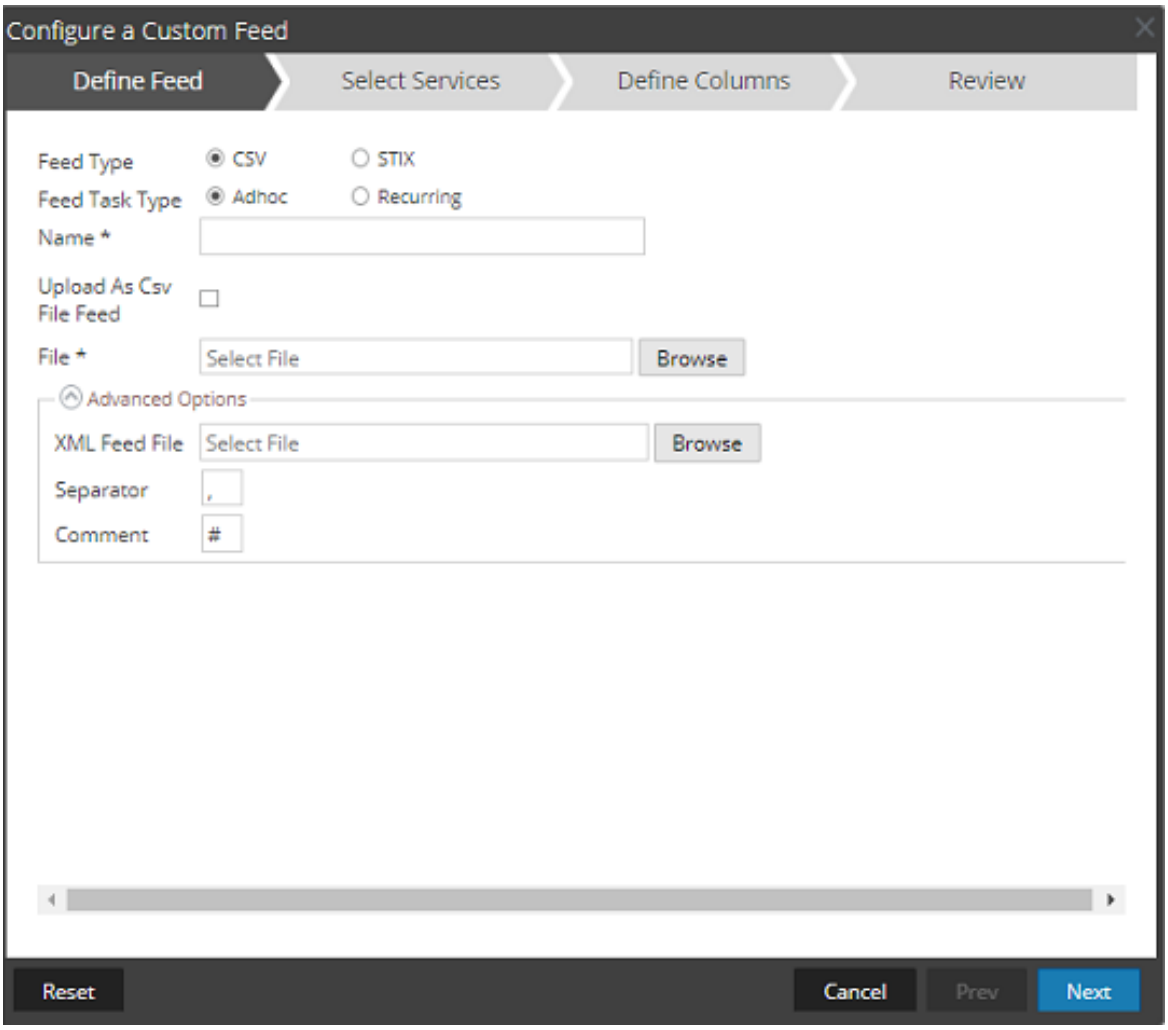
3. Select the Feed Type: **CSV** or **STIX**.
4. To define a feed based on a `.csv` formatted feed data file, select **CSV** (which is the default) in the **Feed Type** field.
5. To define an on-demand feed task that executes once, select **Adhoc** in the **Feed Task Type** field and do one of the following: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

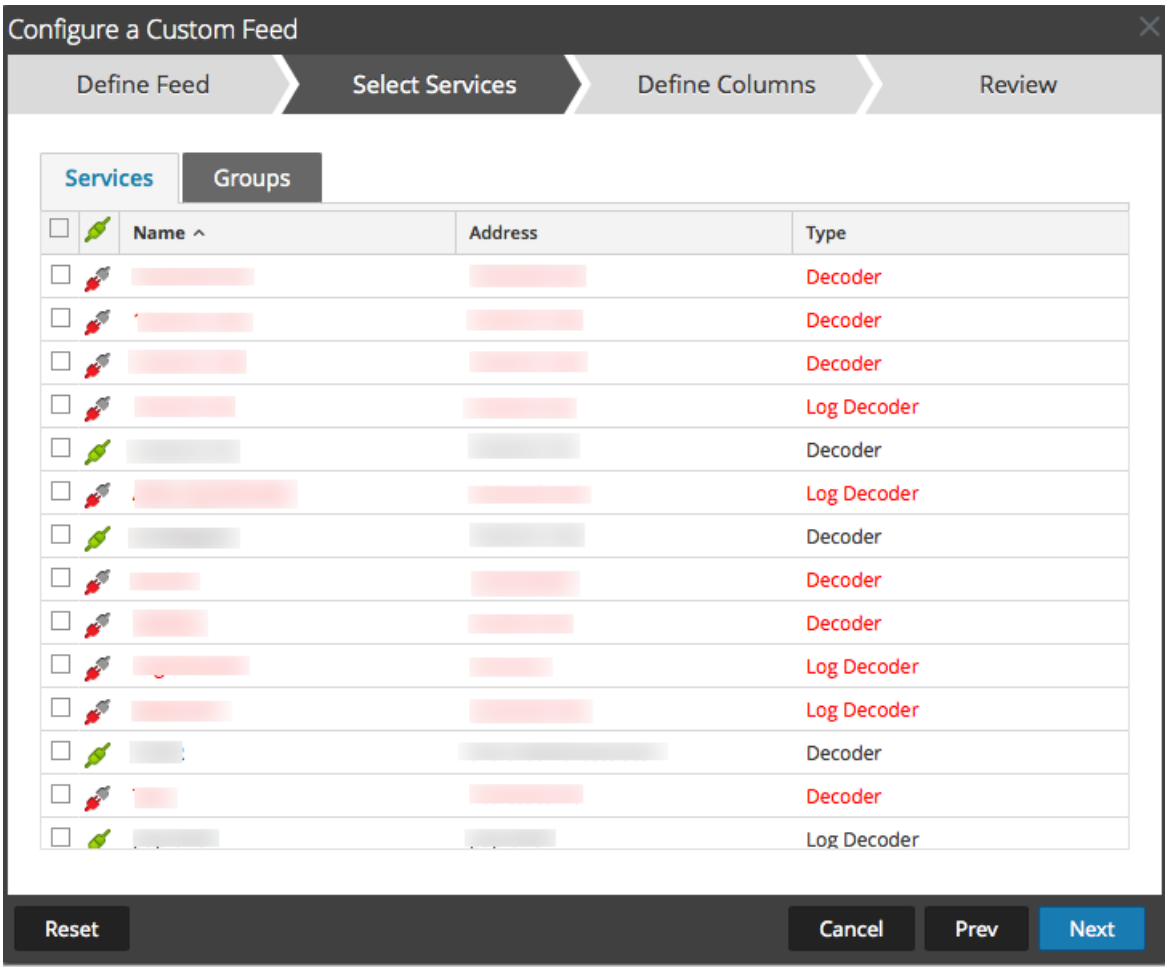
   b. (Conditional) To define a feed based on an XML feed file, select **Advanced Options**.

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

      The Advanced Options are displayed:



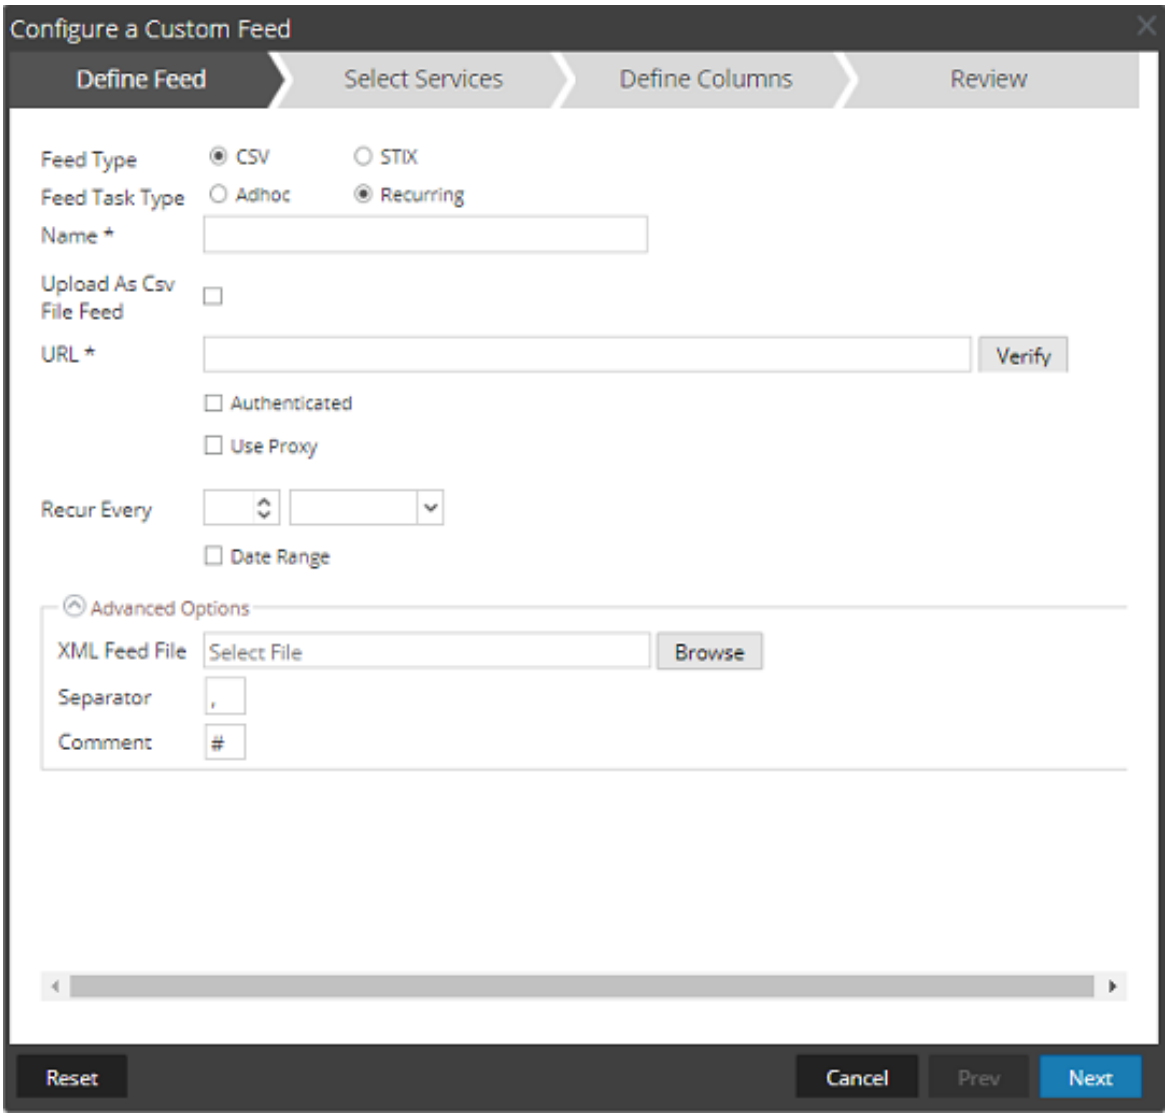
c. Select an XML feed file from the local file system, choose the separator (default is comma), specify the comment characters used in the feed data file (default is #), and click **Next**.
The Select Services form is displayed. This is an example of the form for a feed based on a feed data file with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



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.



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 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 provides your user name and password for authentication to the URL.

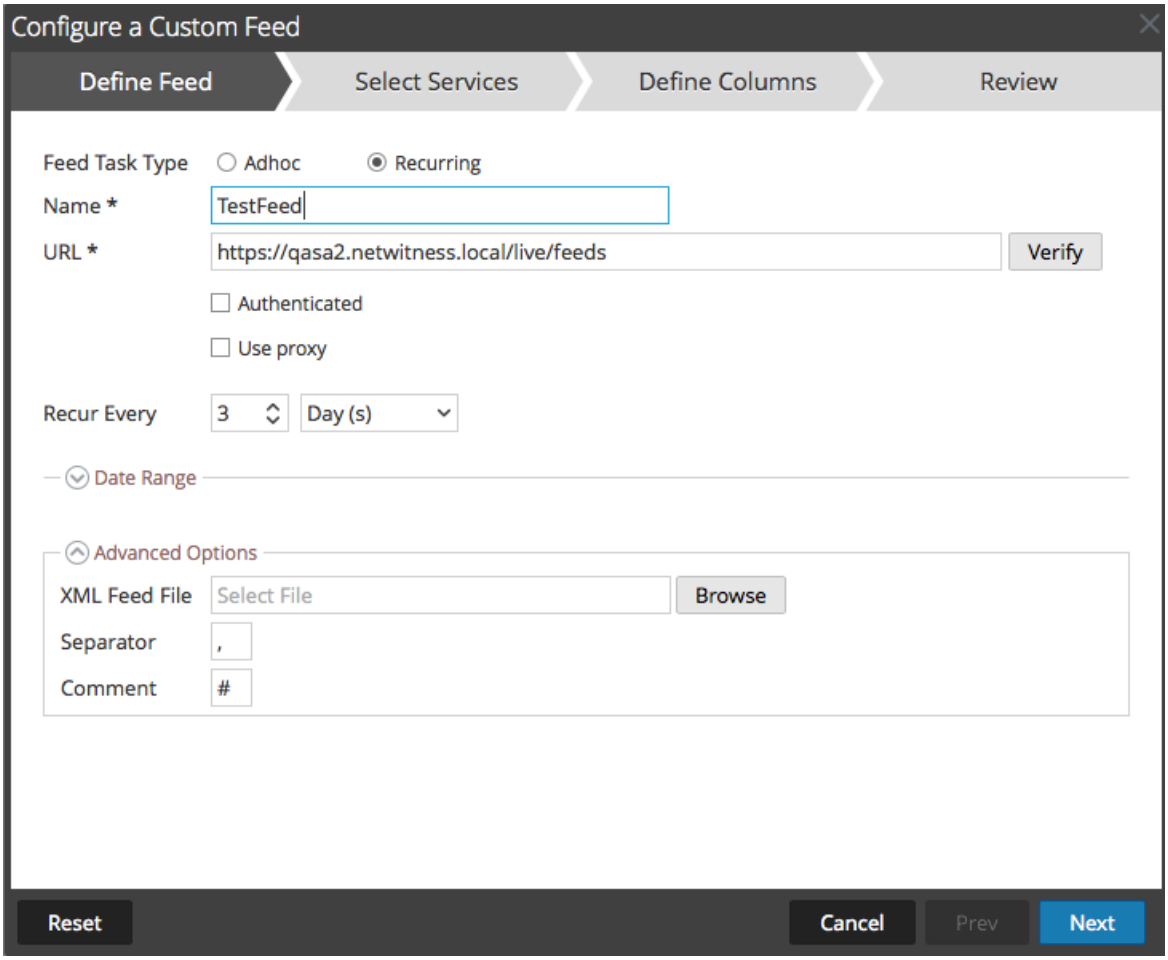
d. If you want the NetWitness server to access the Feed URL through a proxy, select **Use Proxy**. For more information on configuring a proxy, see "Configure Proxy for NetWitness" in the *System Configuration Guide*. By default, the **Use Proxy** checkbox is not set.

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.

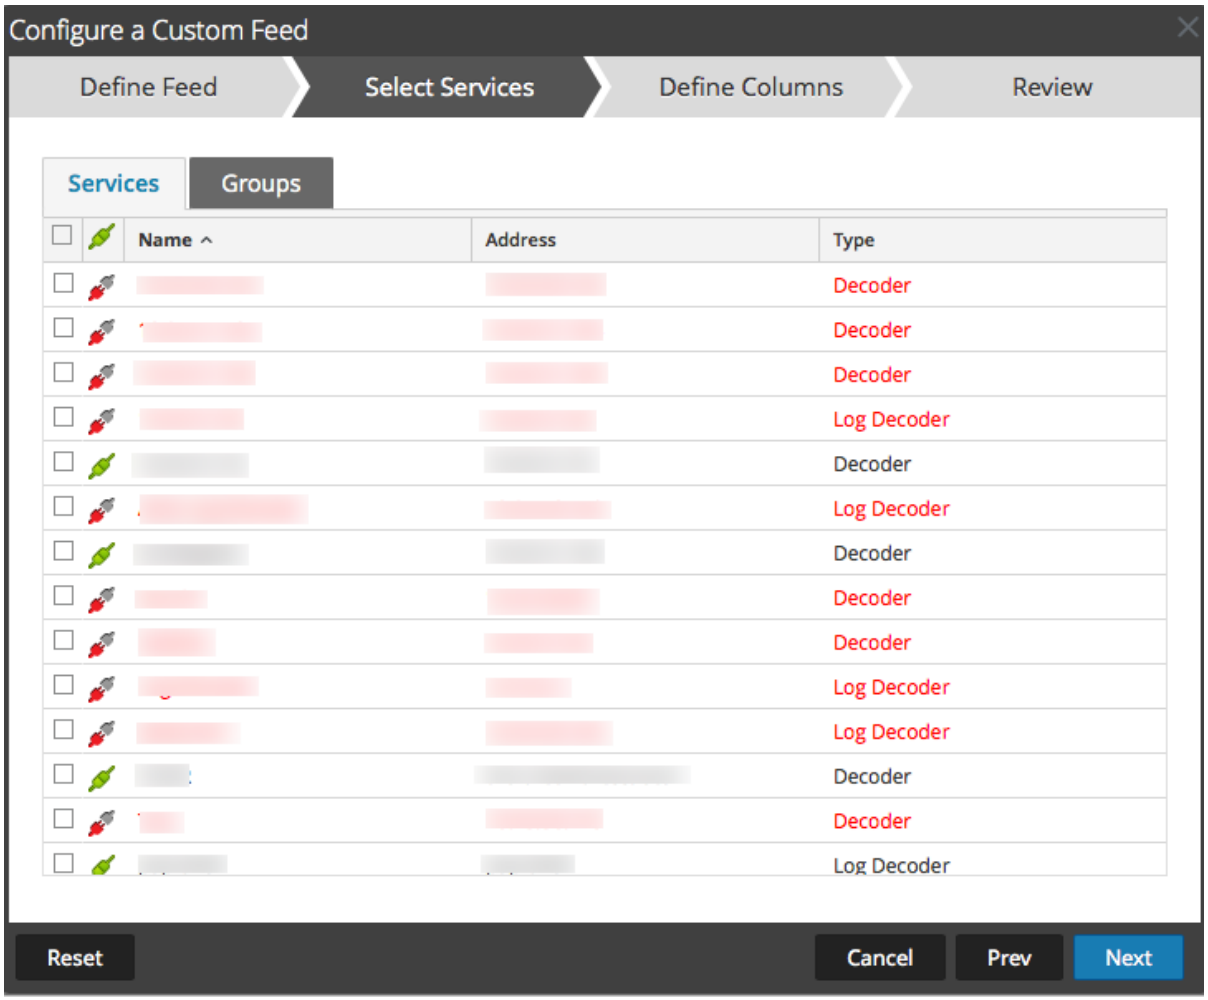
f. To define the date range for the execution of the feed to recur, specify the **Start Date** and time and the **End Date** and time.



7. (Conditional) If you want to define a feed based on an XML feed file:

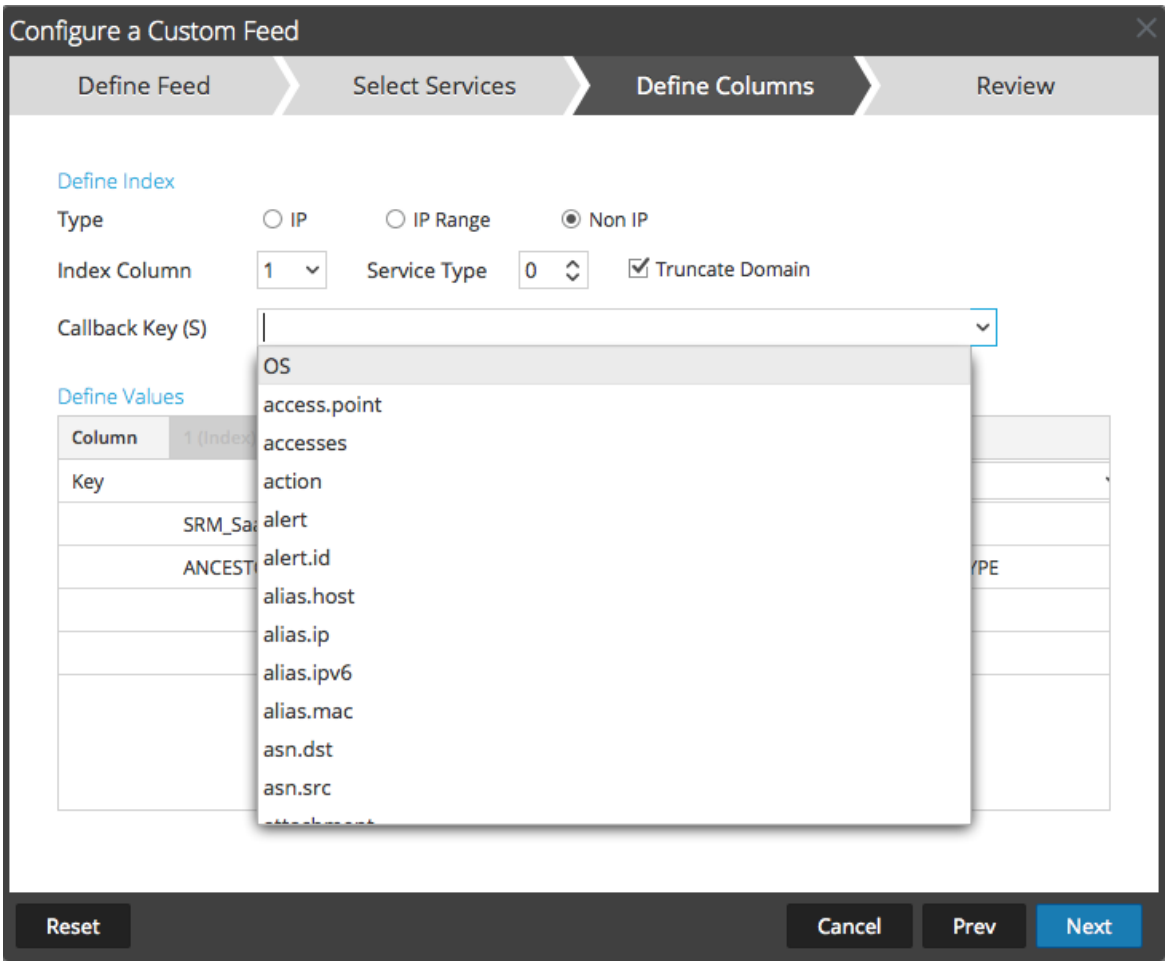
   a. Type the feed **Name**, select **Advanced Options**. The Advanced Options fields are displayed.

   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.

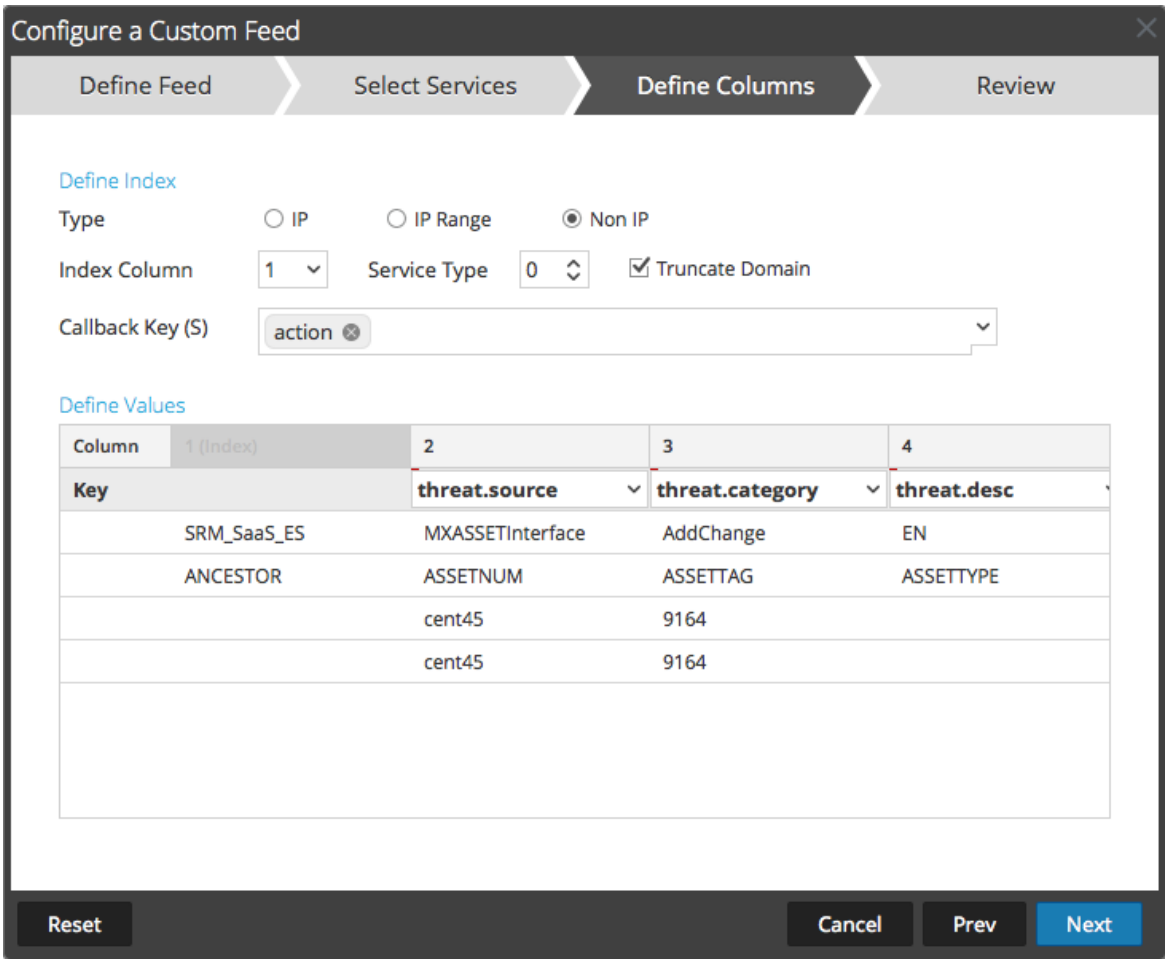


8. To identify services on which to deploy the feed, do one of the following:
   a. Select one or more Decoders and Log Decoders, and click **Next**
   b. Click the **Groups** tab and select a group. Click **Next**.

      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.

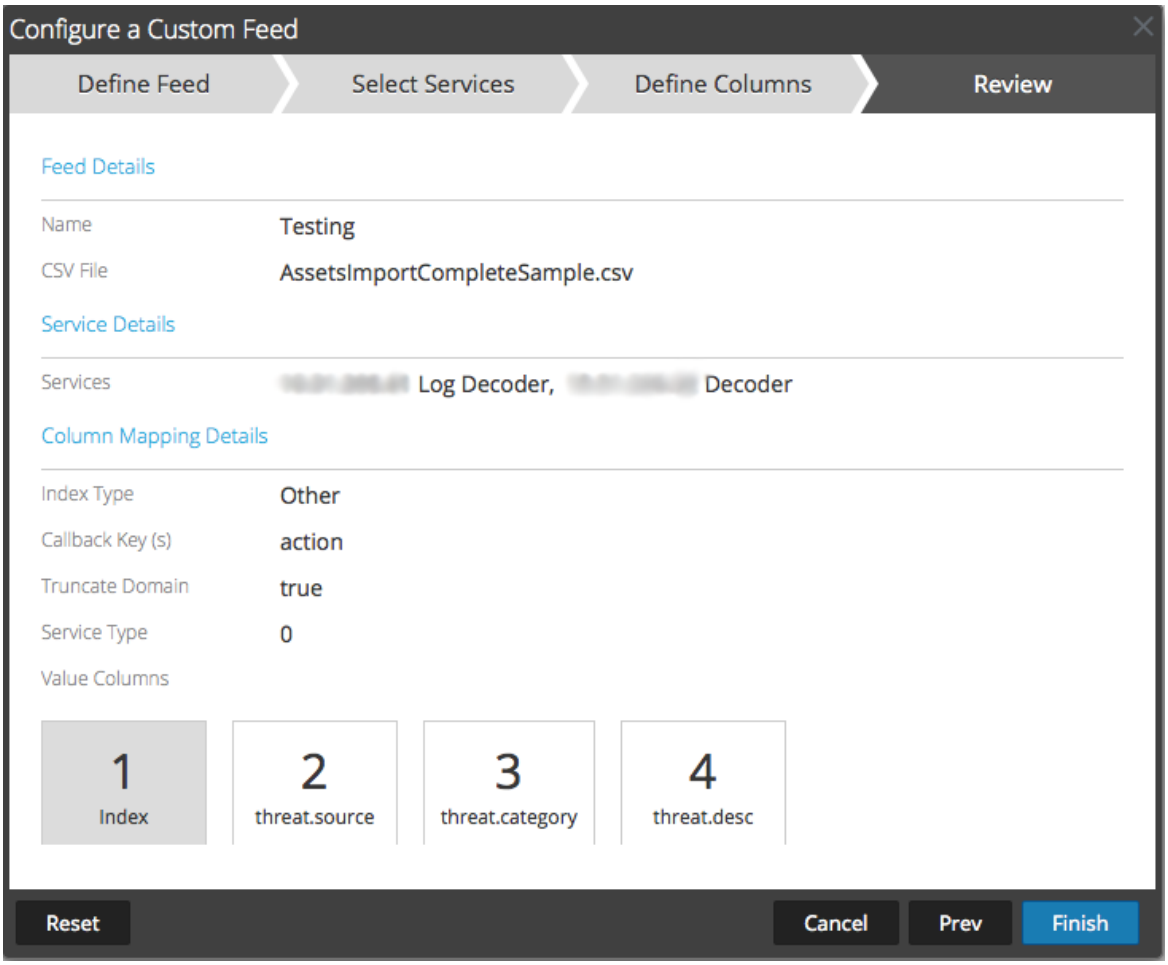


d. Select the language key to apply to the data in each column from the drop-down list. The meta keys displayed in the drop-down list is based on the meta keys available for the service define values. You can also add other meta keys based on advanced expertise.



e. Click **Next**.

The Review form is displayed.



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.

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. NetWitness 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 to NetWitness latest version, you must re-configure the STIX recurring feed.

In NetWitness, 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.

The STIX and TAXII communities work closely together to ensure that they continue to provide a full stack for sharing threat intelligence.

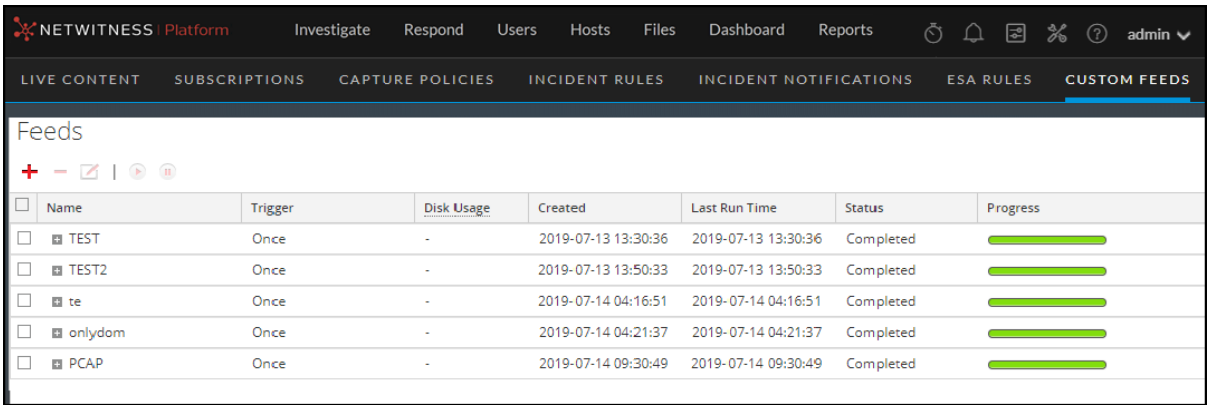
Apart from the TAXII server, STIX data can also reside on a REST server and you can fetch the STIX file from the REST server by providing the URL of the REST server. For example, `http://stixrestserver.internal.com`.

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 server.

To create a STIX custom feed:

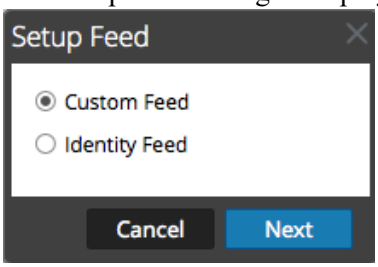
1. Go to **(Configure) > CUSTOM FEEDS**.

   The Custom Feeds view is displayed.



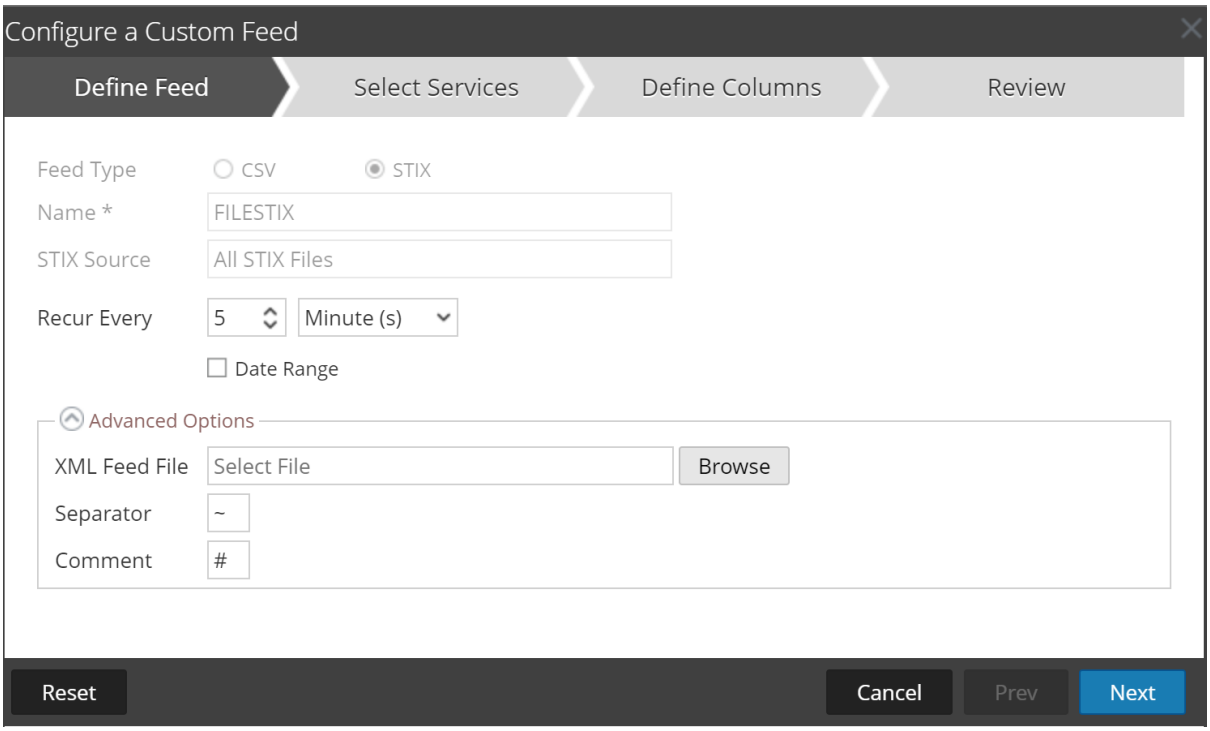
2. In the toolbar, click  .

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.



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 verifies the connection to the server, so that NetWitness 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.

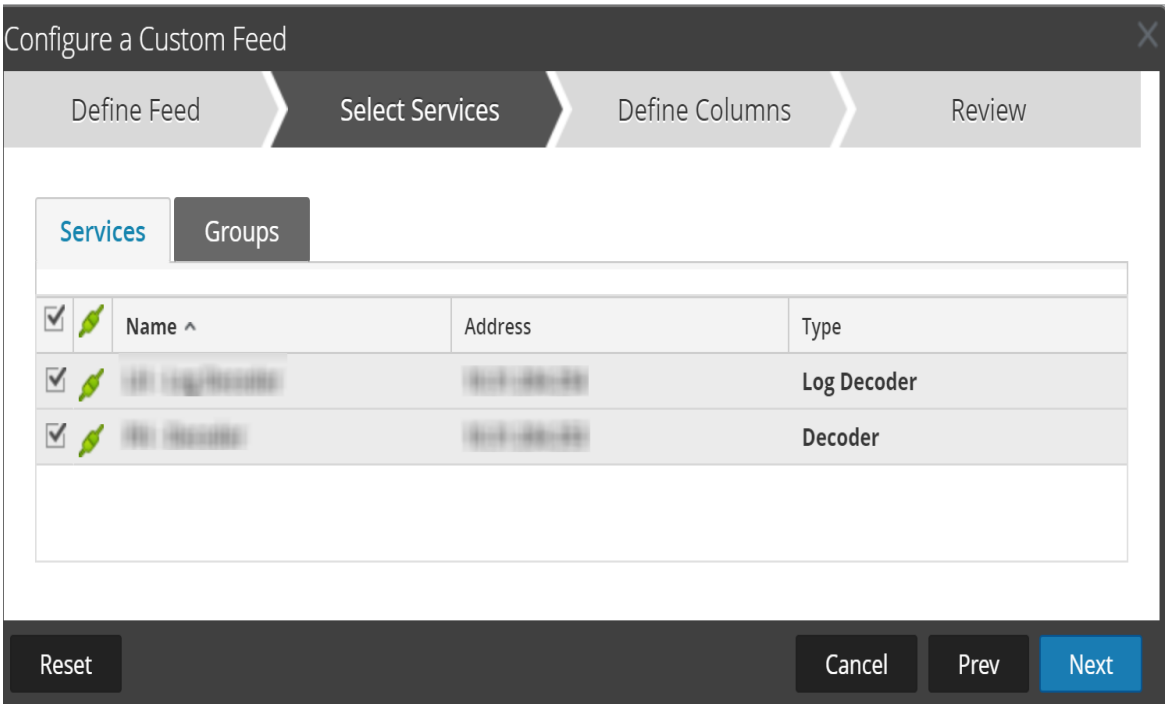
5. (Optional) Select **Advanced Options**,to define a feed based on an XML feed file.

   a. **XML Feed file**: Browse and select an XML feed file from the local file system.

   b. **Separator**: Choose a separator (default is comma).

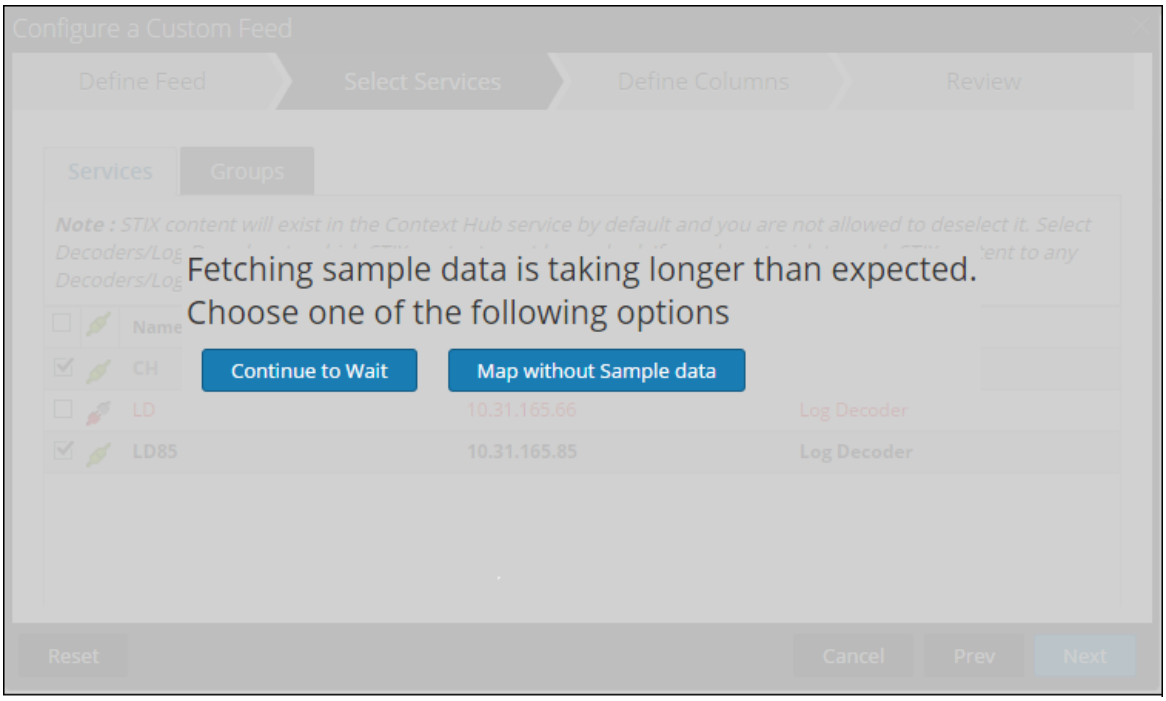
c. **Comment**: Specify the comment characters used in the feed data file (default is #).

6. Click **Next**.
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.



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**.

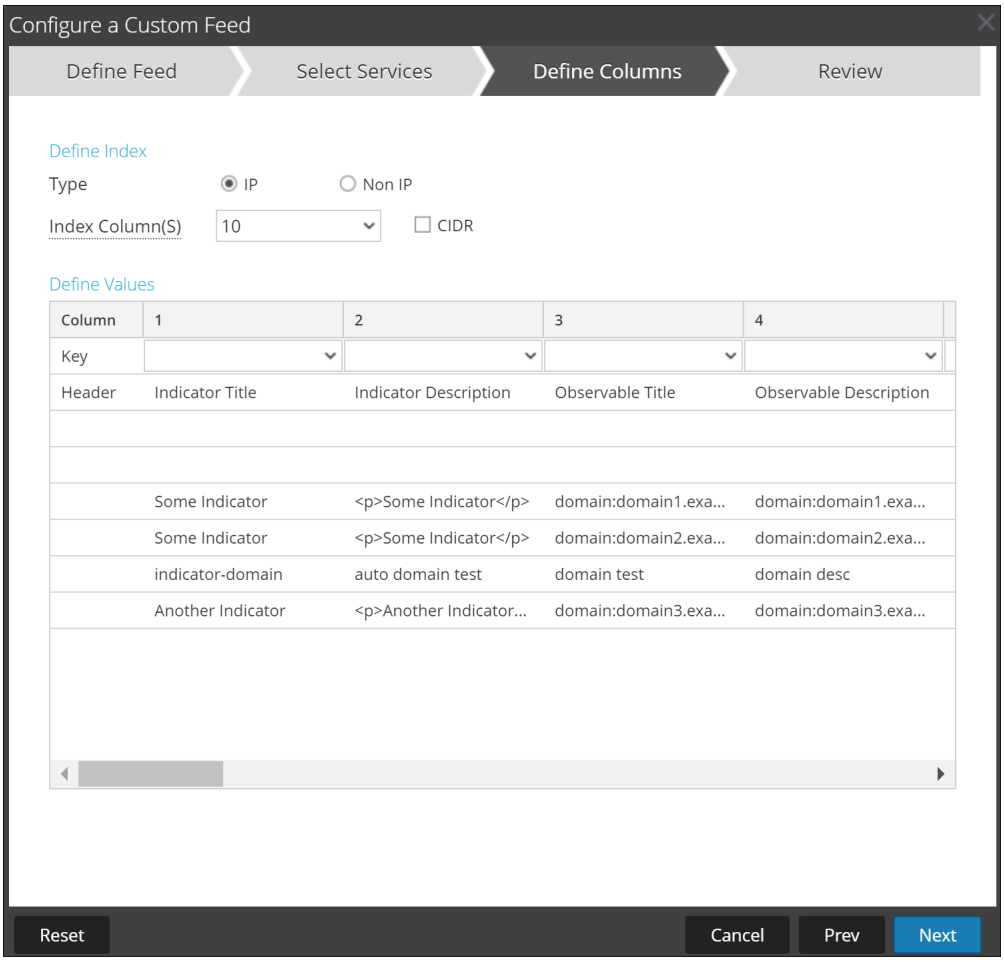
      If the data from the STIX server is large, the following message is displayed:



- 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.

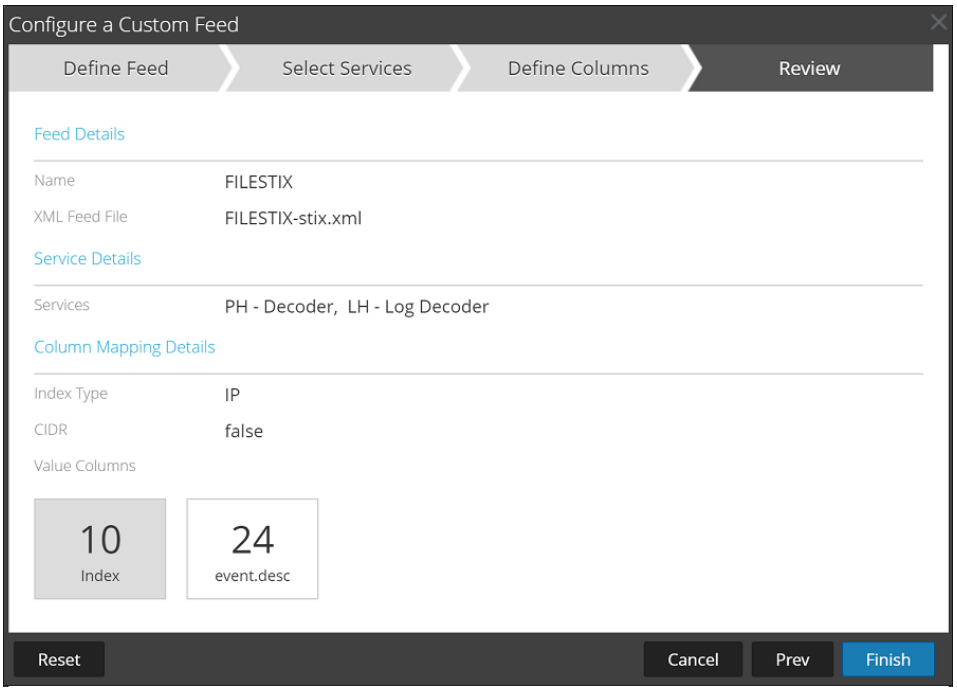


Note:
- If the **Index Type** is Non IP, you can select multiple index columns in the **Index Column(S)**. The values from all the selected columns are merged in the first index column that you selected and the merged values are pushed to the Log Decoder for parsing. For example, in the **Index Column(S)** 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.



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.

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*.

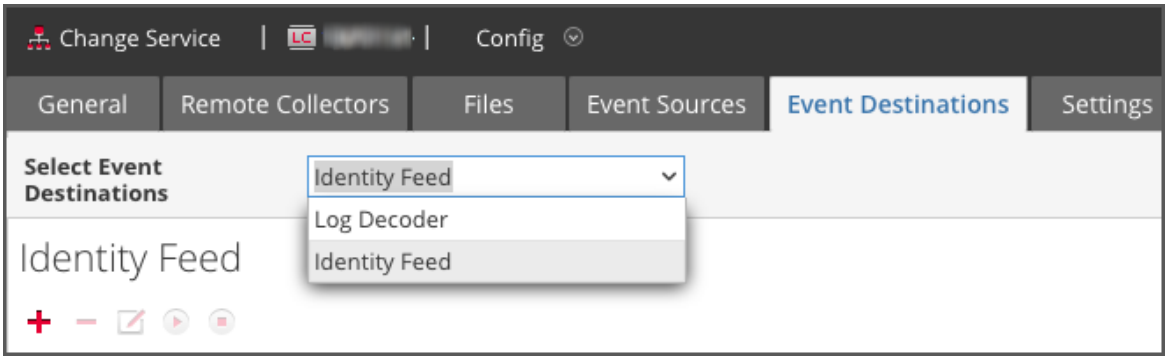
Create an Identity Feed

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

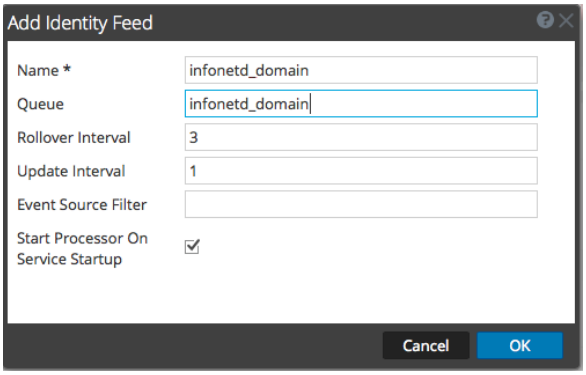
To create an identity feed:

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.

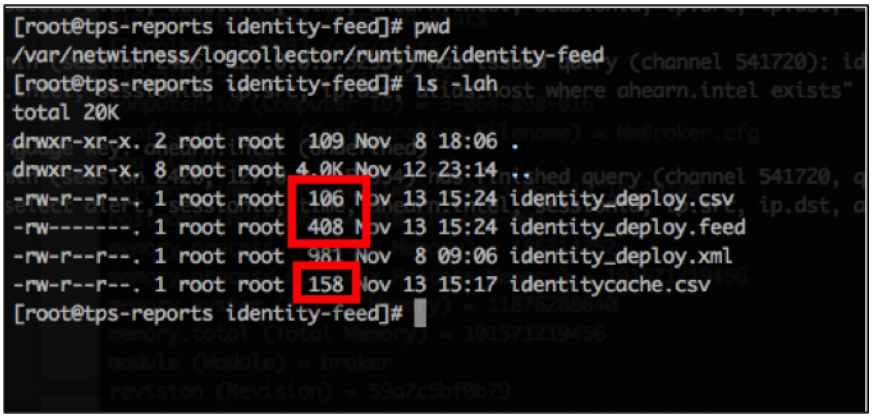
   The Queue name identifies the feed within the Log Collector. Use the name of the feed for the Queue.



   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.

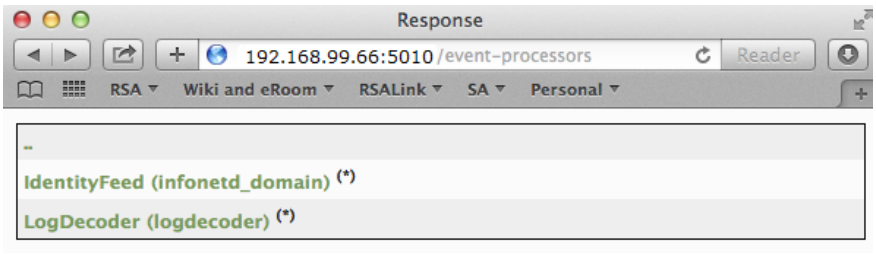
b. Verify that data is written to the feed files. SSH to the Log Decoder/Collector or Virtual Log Collector being configured. Navigate to `/var/netwitness/logcollector/runtime/identity-feed` and verify that the `Identity_deploy` files are getting populated with data.



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:

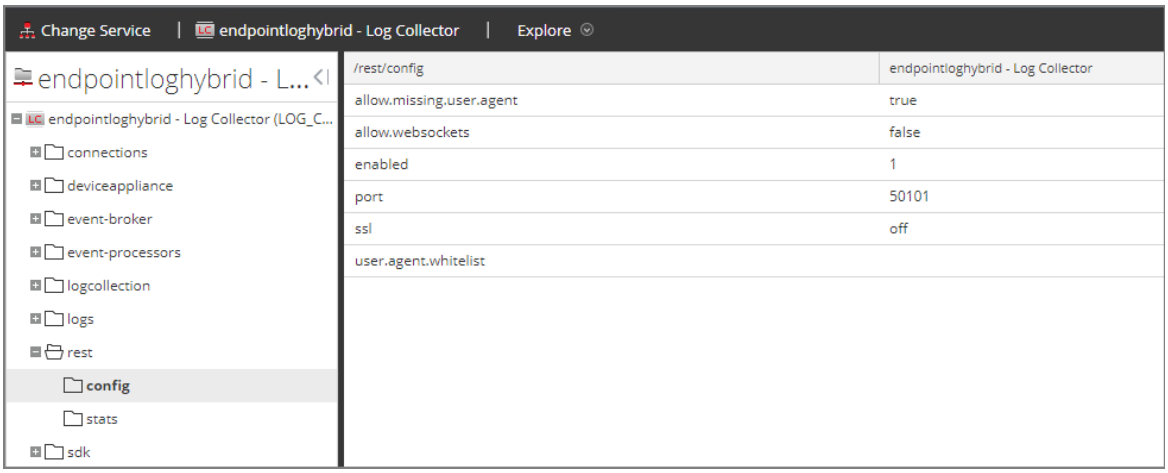


The screen contains the name of the identity feed you created earlier (`infonetd_domain`, in this example).

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**.



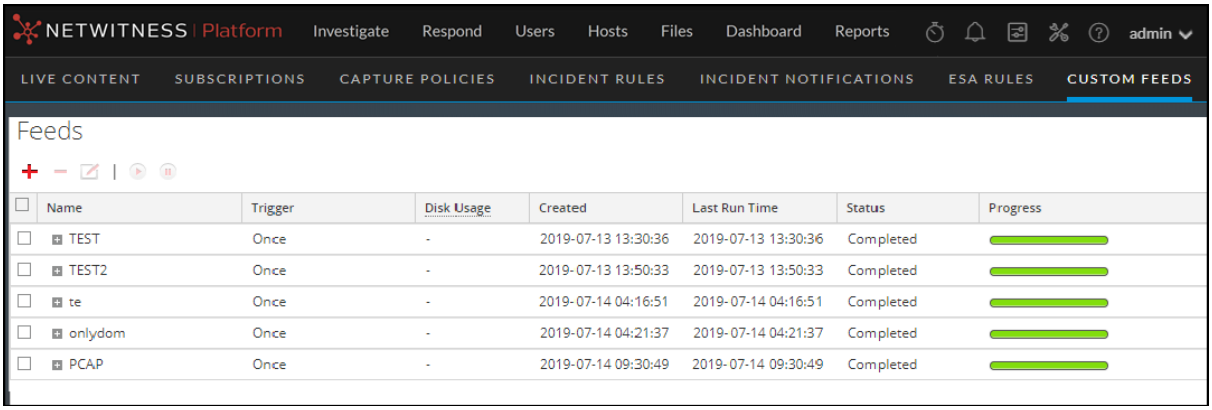
For REST to be active, **enabled** must be set to **1**.

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.

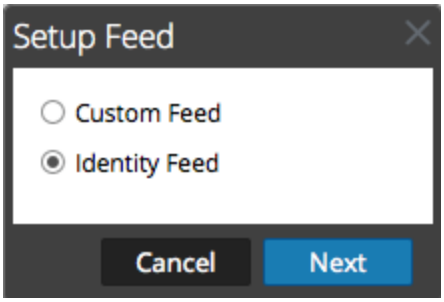
3. Go to **(Configure) > Custom Feeds**.

The Feeds dialog is displayed.



4. In the toolbar, click .

The Setup Feed dialog is displayed.



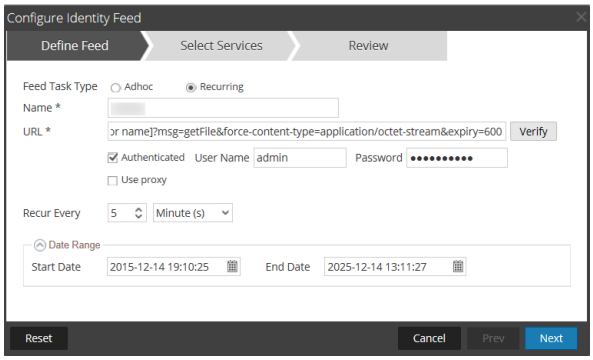
5. Make sure **Identity Feed** is selected and click **Next**.

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: NetWitness verifies the location where the file is stored, so that NetWitness 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 UI server can access the Log Collector's REST API port (50101). This can be tested by going to the NetWitness 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 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 NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

9. To define the recurrence interval, do one of the following:

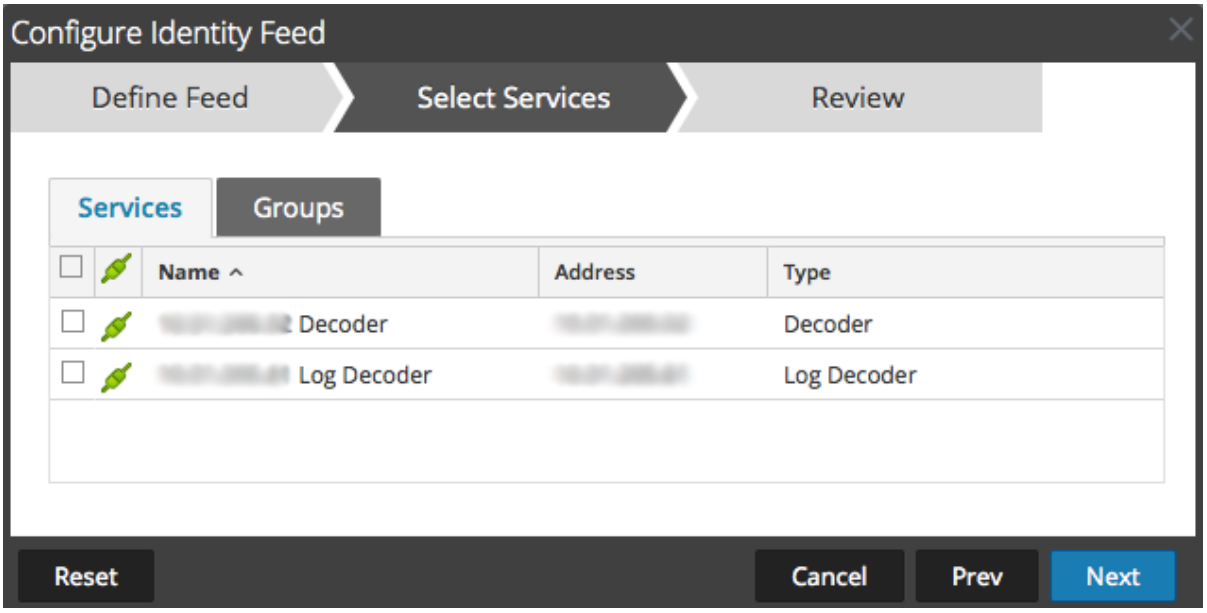
- Specify the number of minutes, hours, or days between recurrences of the feed.
- Enter the date range for the execution of the feed to recur, specify the **Start Date** and time and the **End Date** and time.

10. If using SSL encryption, you need to install the REST API SSL certificate for the Log Collector into the NetWitness 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.

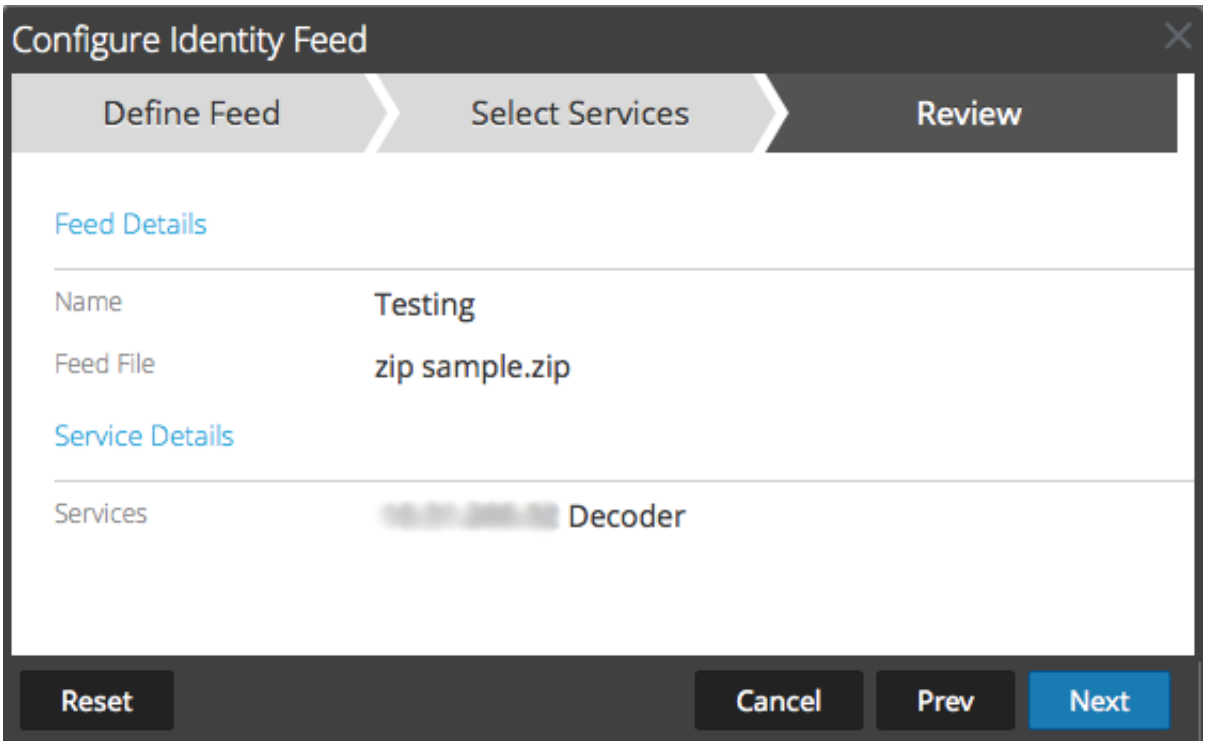
11. Click **Verify** to verify your identity feed configuration before you proceed to the Select Services dialog.
12. Click **Next**.

    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**.

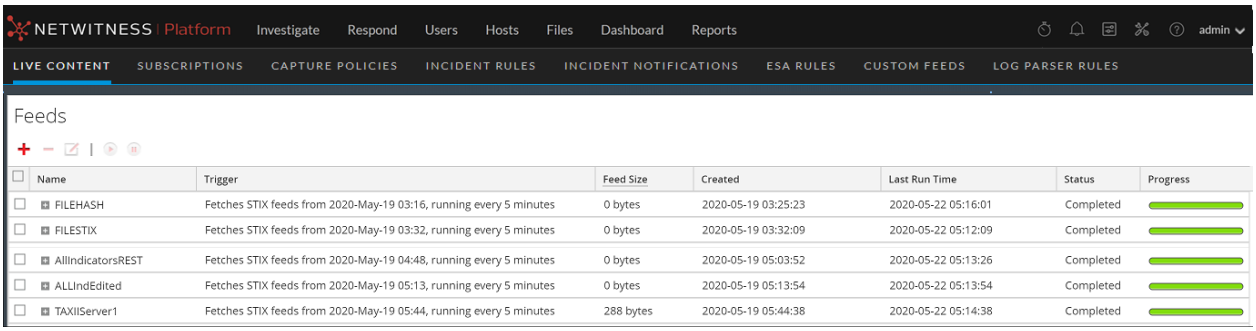
    The Review dialog is displayed.



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.



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 UI server key store. If this certificate is not imported, the NetWitness 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 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 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 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.

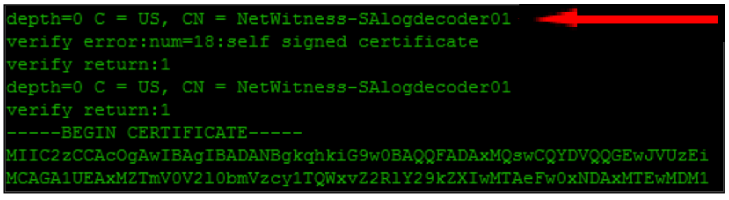
1. SSH to the NetWitness UI server.
2. Run the following command to output the CN name of the SSL cert:

   ```
   echo -n | openssl s_client -connect <log decoder>:50101 | sed -ne '/BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p'
   ```

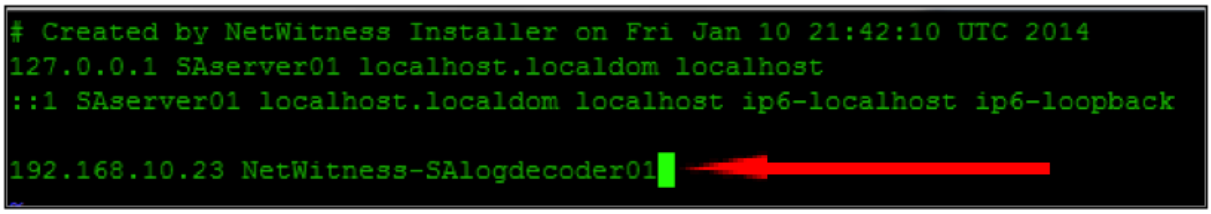
   For example:

   ```
   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.

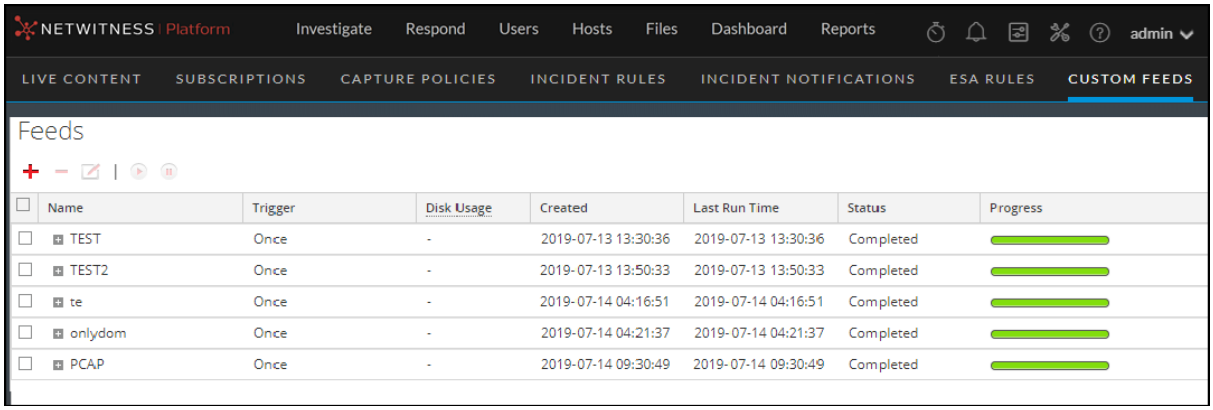
Edit, Upload, or Remove a Feed

You can upload a feed, edit an existing feed, or remove a feed.

To edit an existing feed:

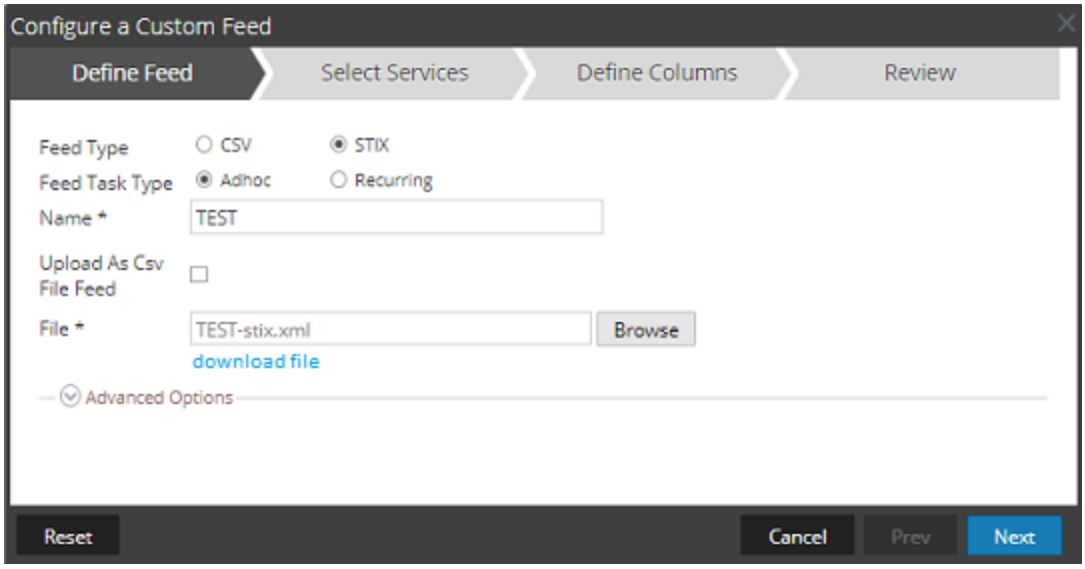
1. Go to (Configure) > **Custom Feeds**.

   The Feeds view is displayed.



2. In the toolbar, select a feed and click .

   The Configure Custom Feed or Configure Identity Feed panel opens in the Custom Feed wizard.



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**.

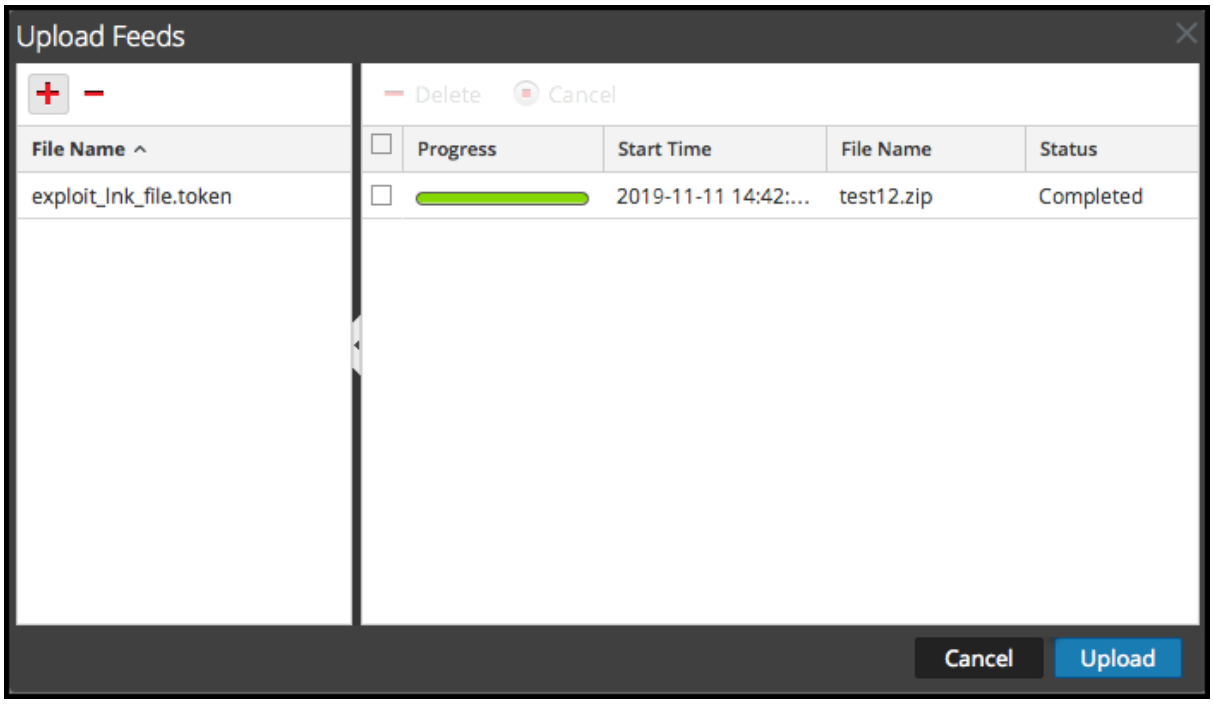
The feed is recreated with the updated file and new feed specifications. The feed is added to the Feeds list and progress bar tracks completion. Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file is listed in the Feeds list. 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  .

   The Upload Feeds dialog is displayed.



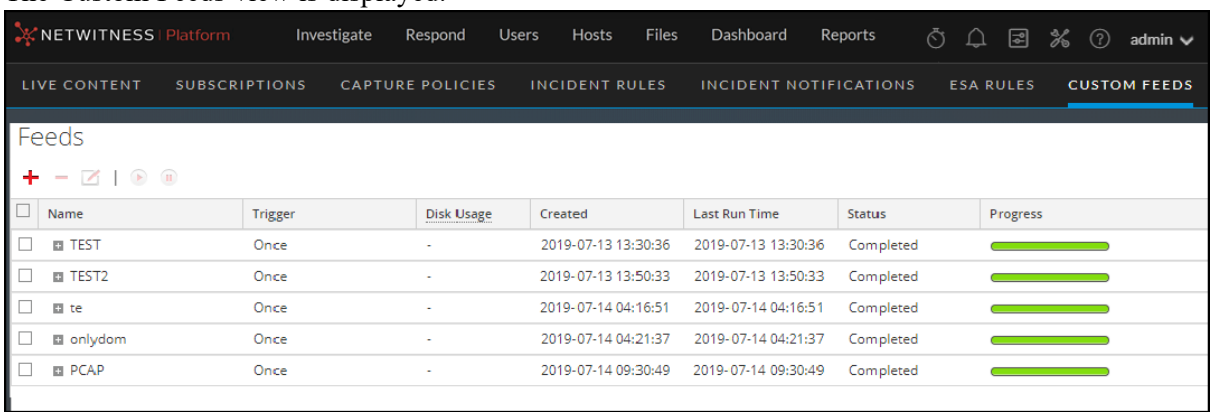
5. In the **File** grid, click + and select a feed file. Supported files are *.feed, *.token, and *.filter.
6. Select the feed file from the **File** list and click **Upload**.

   The Upload Job list is updated to show the progress and status of the uploaded feed.

To remove a feed:

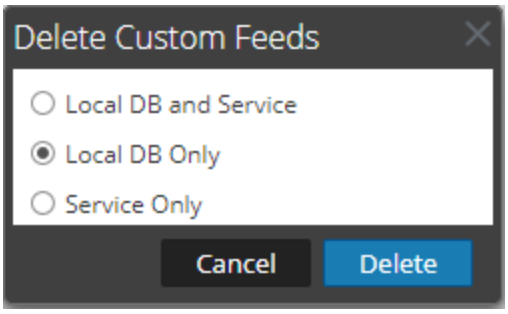
1. Go to (Configure) > **Custom Feeds**.

   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:

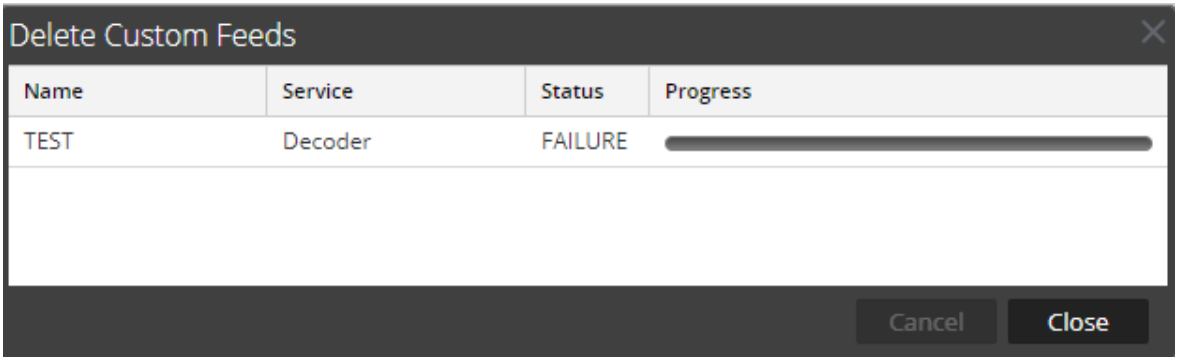
- If you choose to delete the feed from **Local DB and Service**, the feed is deleted from both the service and the local NetWitness box. The deleted feed will no longer be seen on the NetWitness user interface.
- If you choose to delete the feed from **Local DB Only**, the feed is deleted from the local NetWitness box. The deleted feed will not be seen on the NetWitness 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 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 reserved keywords. Do not use these keywords as custom meta keys as it can cause display error in the NetWitness 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

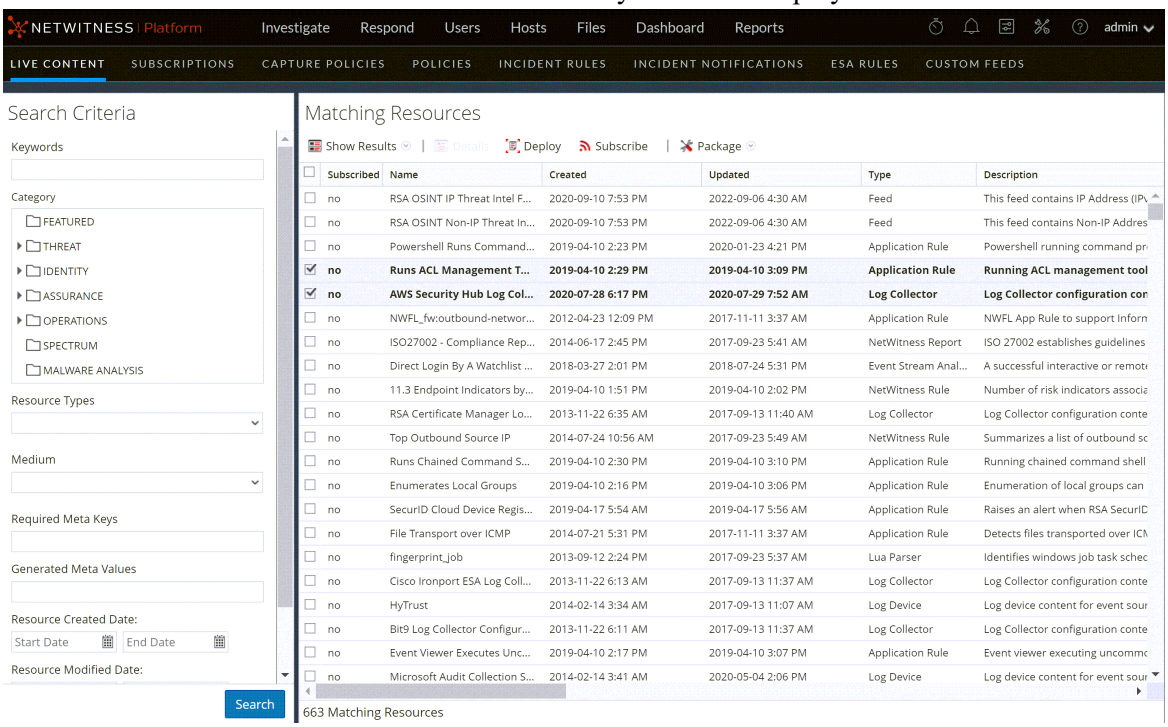
```
score
cards
intro-subject
context-intro
im-risk-score-type
im-priority-type
im-score-type
incident-details
cs-name
im-score-critical
im-score-high
im-score-medium
im-score-low
im-score-lowest
im-score-neutral
alert-score-high
ecat-score
loading-msg
content-loading
live-title
live-text
```

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**.

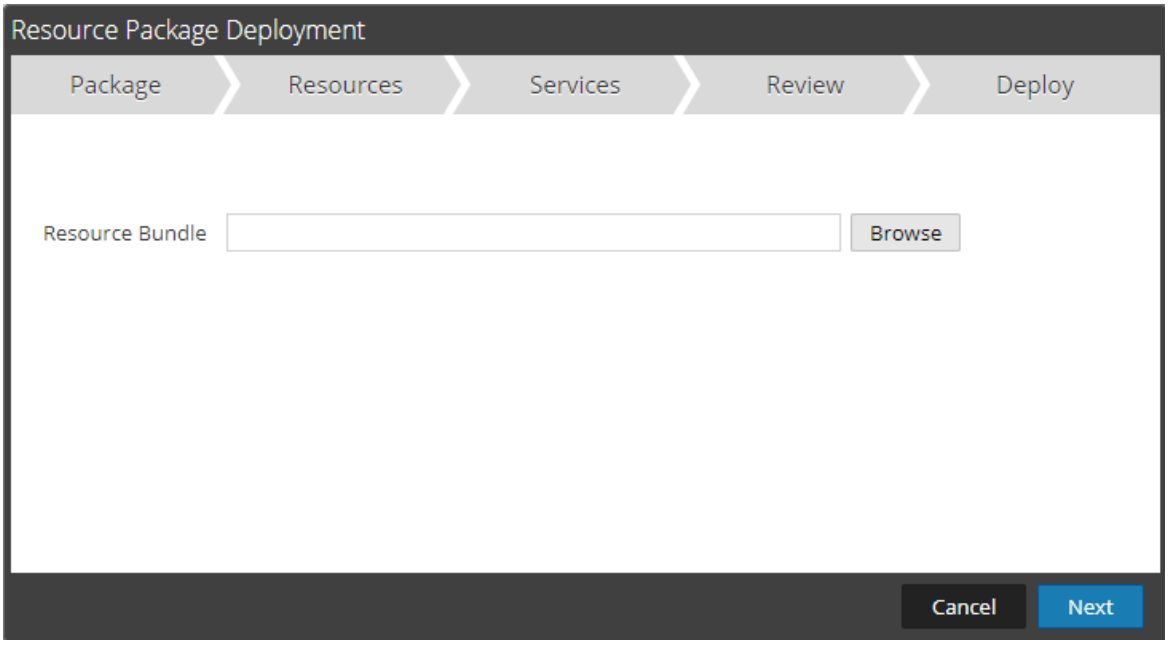
c. Select the checkbox to the left of the resources that you want to deploy.



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.



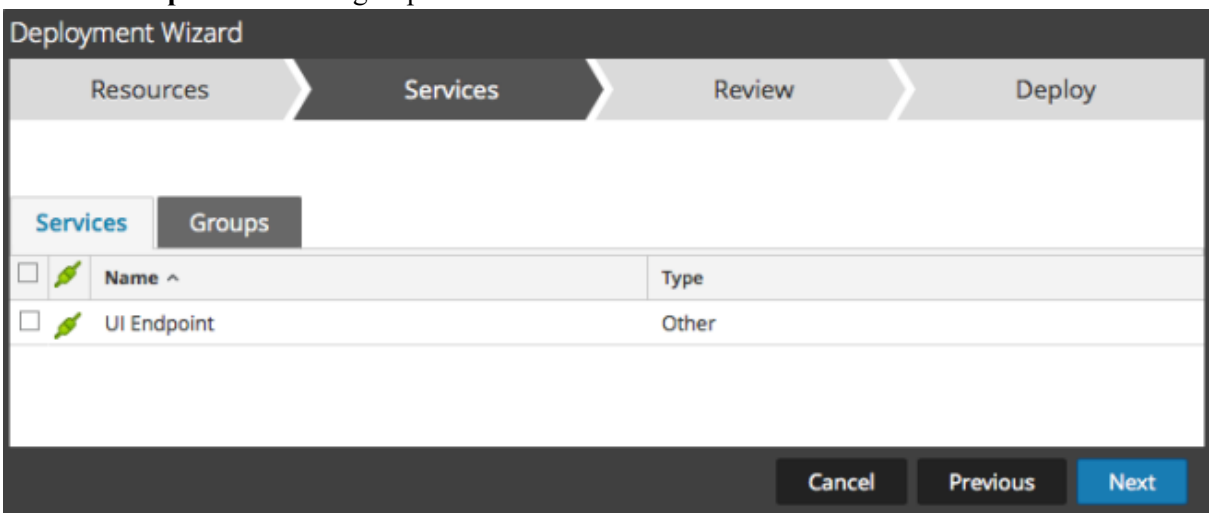
b. Click **Browse** and select a package from your network.

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.

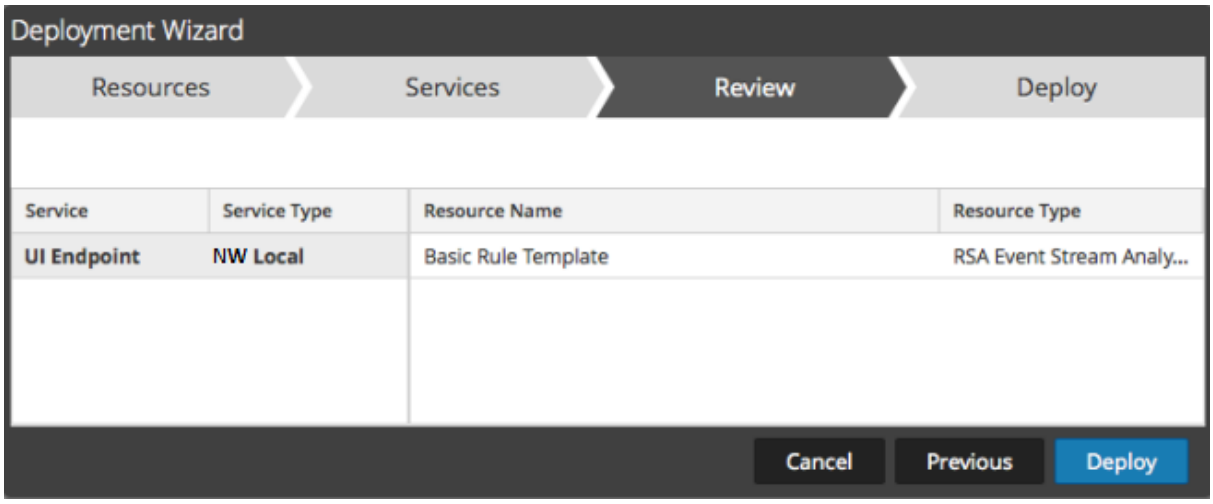
3. Click **Next**.
The **Services** page is displayed. It has two tabs, **Services** and **Groups**, which provide a list of services and service groups that are configured in the (Admin) > Services view. The columns are a subset of the columns available in the Services view.

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.

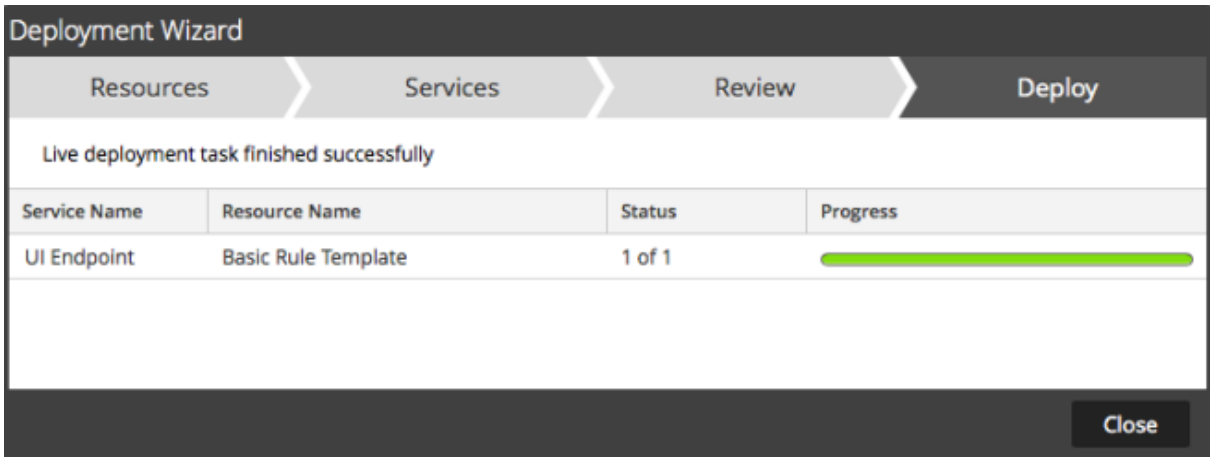


5. Click **Next**.
The **Review** page is displayed.

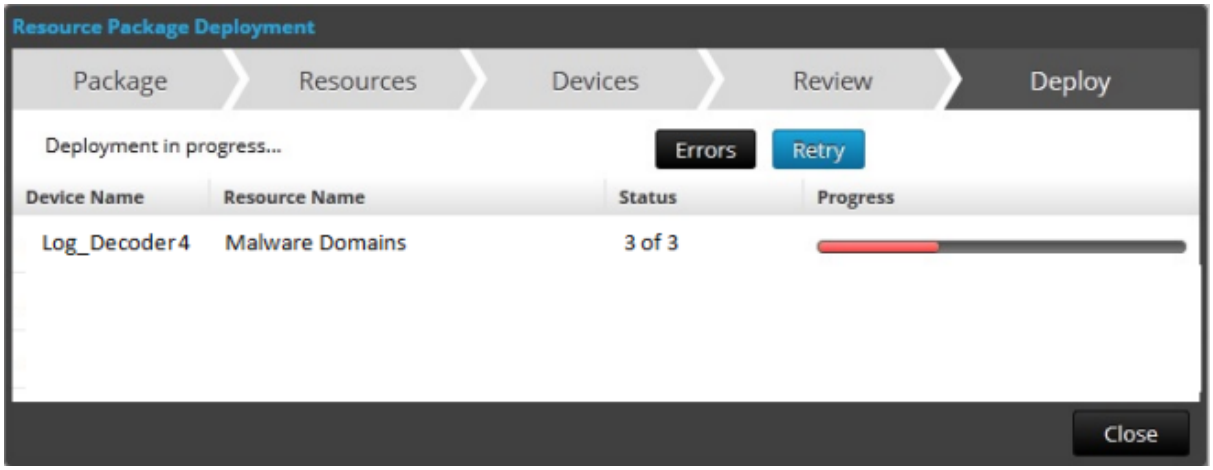


Make sure that you have selected correct resources and the services to which you want to deploy them.

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 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.

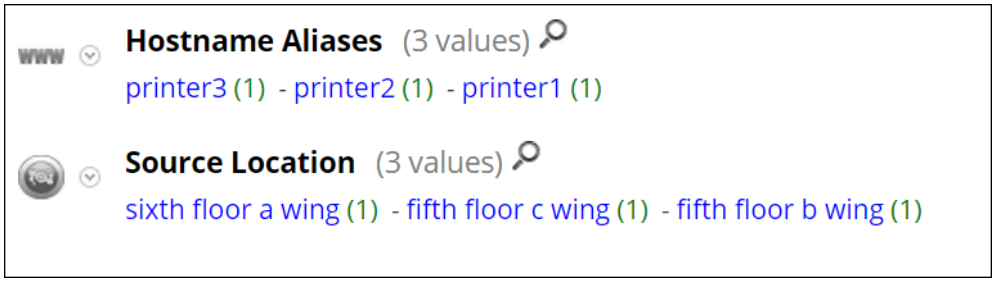
Investigate a Custom Meta Key

> **Note:** You have to log out and log back in to the NetWitness 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**.



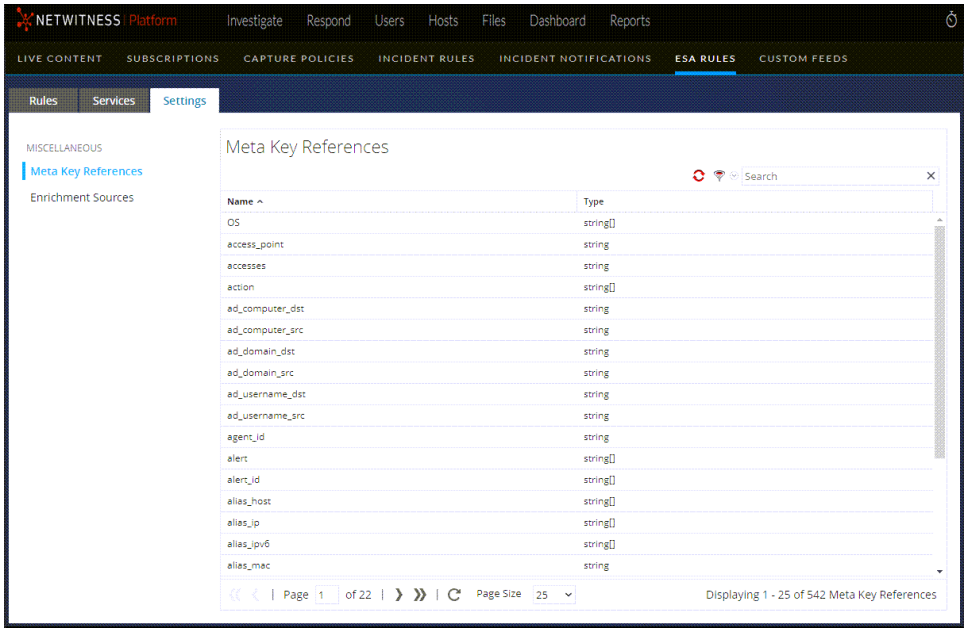
Additional Procedures

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.

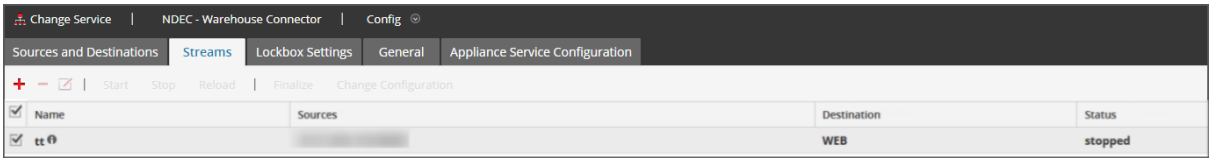
Update the Schema in Warehouse Connector

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

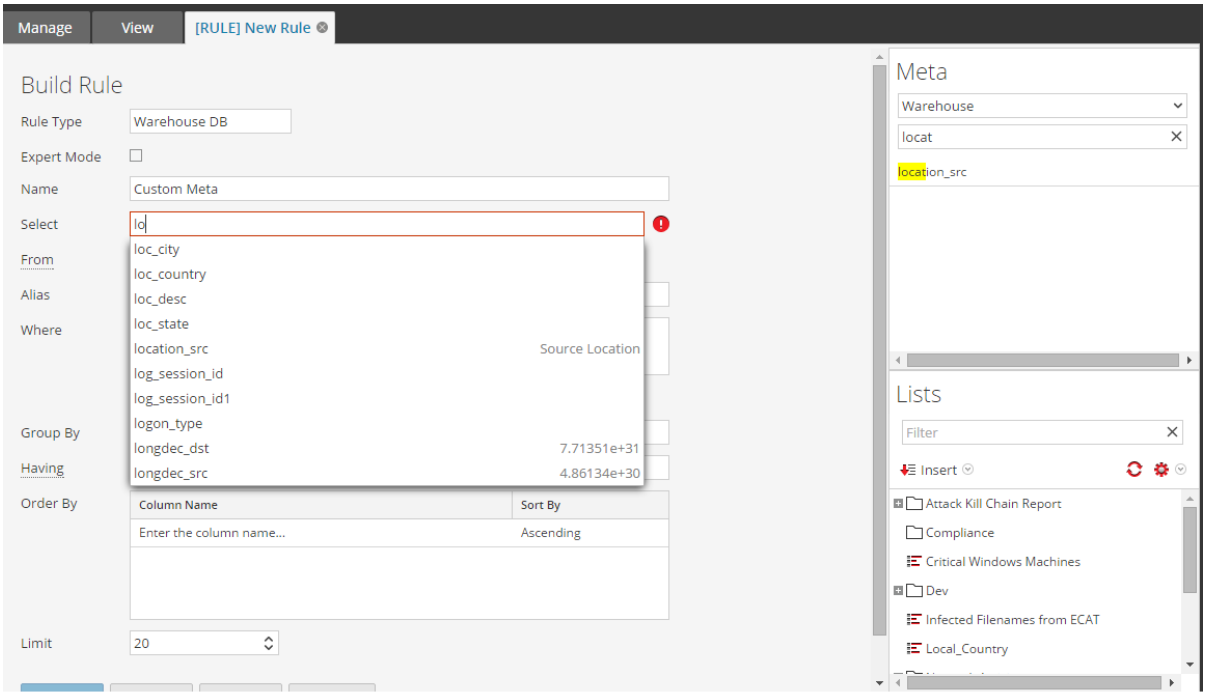
To update the schema in Reporting Engine:

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.



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
- 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 (TLS 1.2)
- Decrypt Incoming Packets TLS 1.3
- 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 Log Decoder
- Upload a Log File to a Log Decoder
- Upload a Packet Capture File
- Troubleshooting Packet Drops

Configure High Speed Packet Capture Capability

This topic guides administrators on how to tune a Network Decoder specifically for high speed packet capture using NetWitness. 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.

NetWitness 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: 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 6 Decoder
- An Intel 82599-based ethernet card, such as the Intel x520 or an Intel i40e-based card like the Intel XL710. All NetWitness-provided 10G cards meet this requirement such as the following:
  - All SMC-10GE cards provided by NetWitness.
  - A Dell Network Daughter Card using an Intel controller to provide 10G network interfaces.
- 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

NetWitness Platform 12.3 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 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/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 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 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

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 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.

IMPORTANT: If you perform the database stagger operation from the UI Explore page, it takes a long time to complete the operation based on the data and results in a timeout. NetWitness recommends that you use the RESTful API or NwConsole to perform the database stagger operation.

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, nwmalwaredomains, 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 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 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 High Speed Packet Capture Capability

This topic guides administrators on how to tune a Network Decoder specifically for high speed packet capture using NetWitness. 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.

NetWitness 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: 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 6 Decoder
- An Intel 82599-based ethernet card, such as the Intel x520 or an Intel i40e-based card like the Intel XL710. All NetWitness-provided 10G cards meet this requirement such as the following:
  - All SMC-10GE cards provided by NetWitness.
  - A Dell Network Daughter Card using an Intel controller to provide 10G network interfaces.
- 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

NetWitness Platform 12.3 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 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/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 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 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

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 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.

IMPORTANT: If you perform the database stagger operation from the UI Explore page, it takes a long time to complete the operation based on the data and results in a timeout. NetWitness recommends that you use the RESTful API or NwConsole to perform the database stagger operation.

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, nwmalwaredomains, 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 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 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:

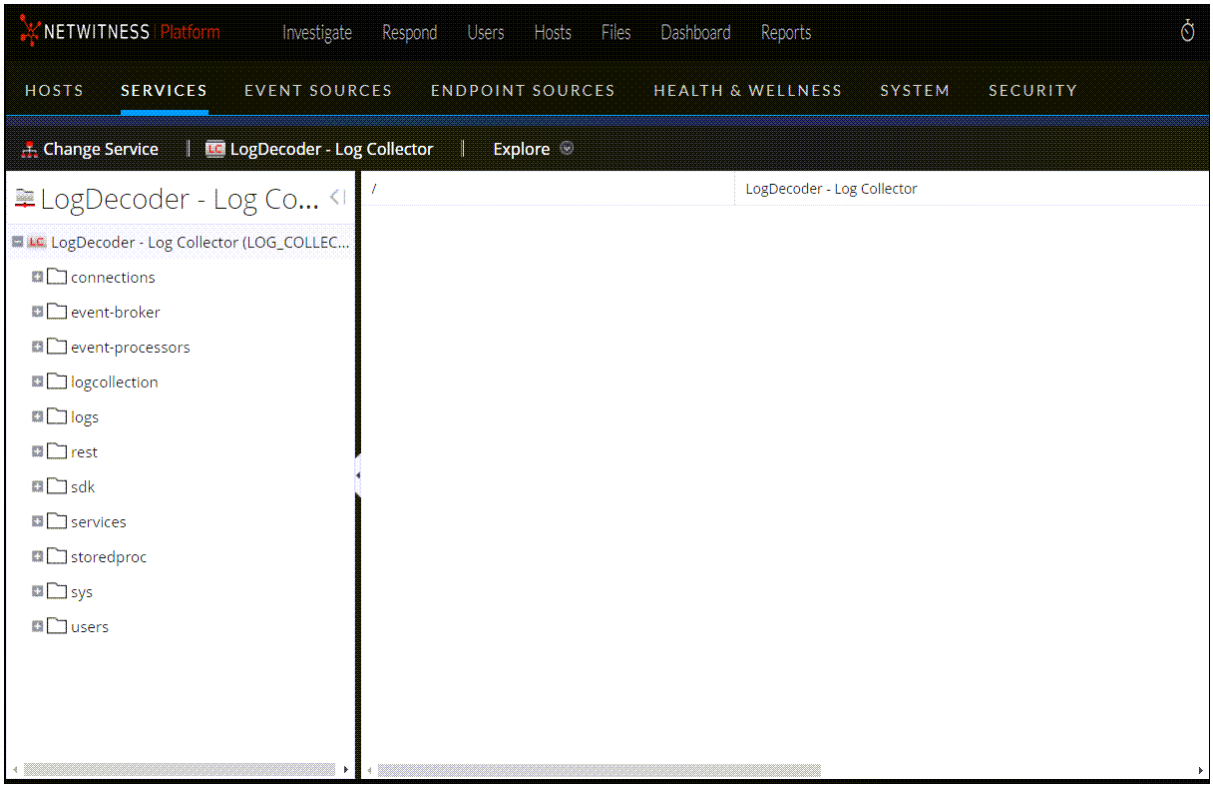
- Disable the Snort rule related to SMB traffic (for example, `M.HackTool.SMB.Impacket-Obfuscation.[Service Names]`). It is recommended because large regex parameters can cause Snort parser to utilize more than 50% CPU.
- Some native parsers should be disabled in some environments. For example, the DNS 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.

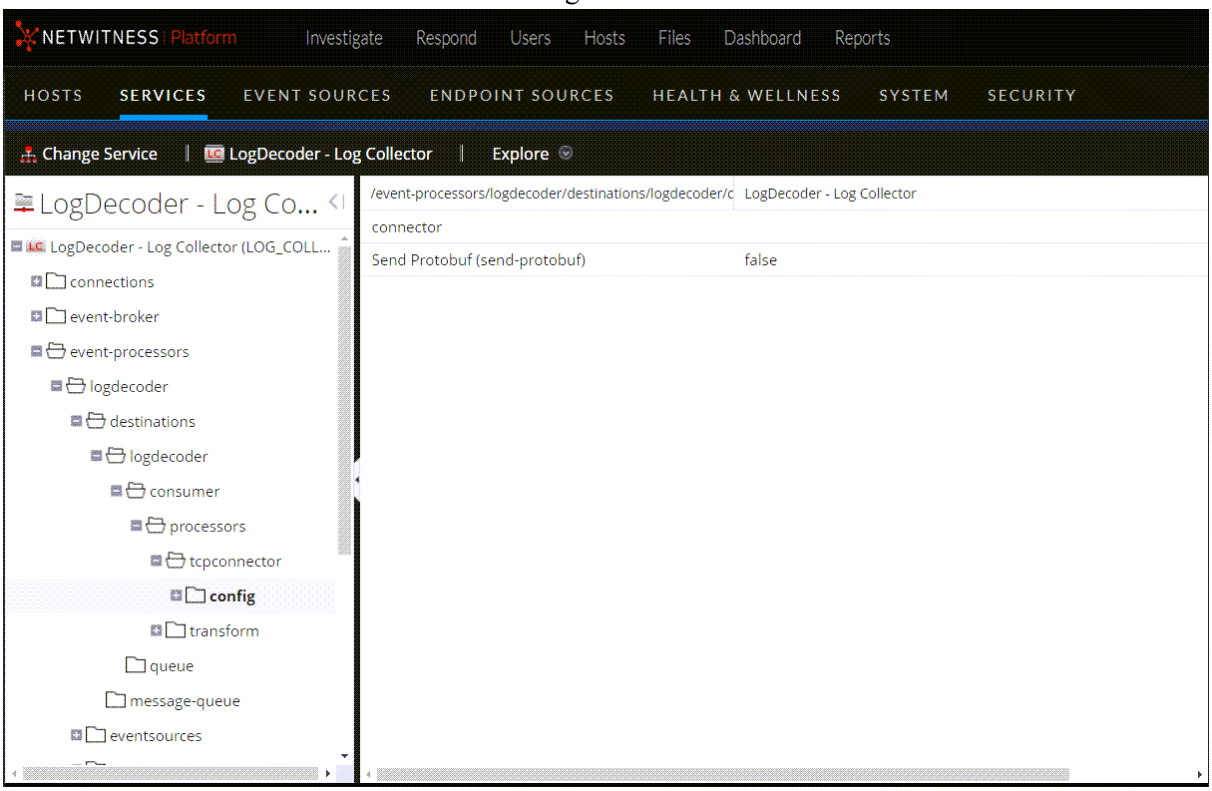
To import a log file to a Log Decoder:

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.size.max`) or has been inactive for a period of time, the session is split. NetWitness 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/assember.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 `sdk` 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.

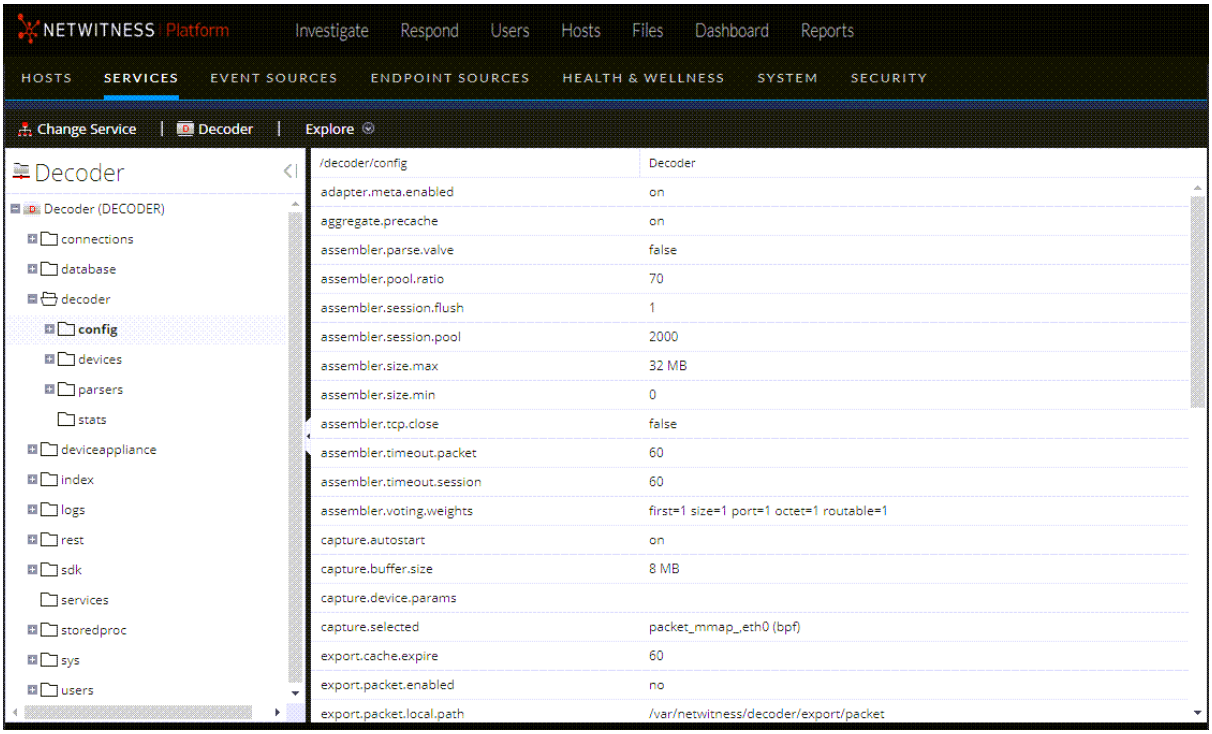
If both timeouts are disabled, the sessions are still split due to time or size expiration. However, the Decoder keeps track of the network stream for as long as it has sufficient memory. Thus, when more packets arrive on the same network stream, the Decoder adds `split` meta items to the subsequent 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.

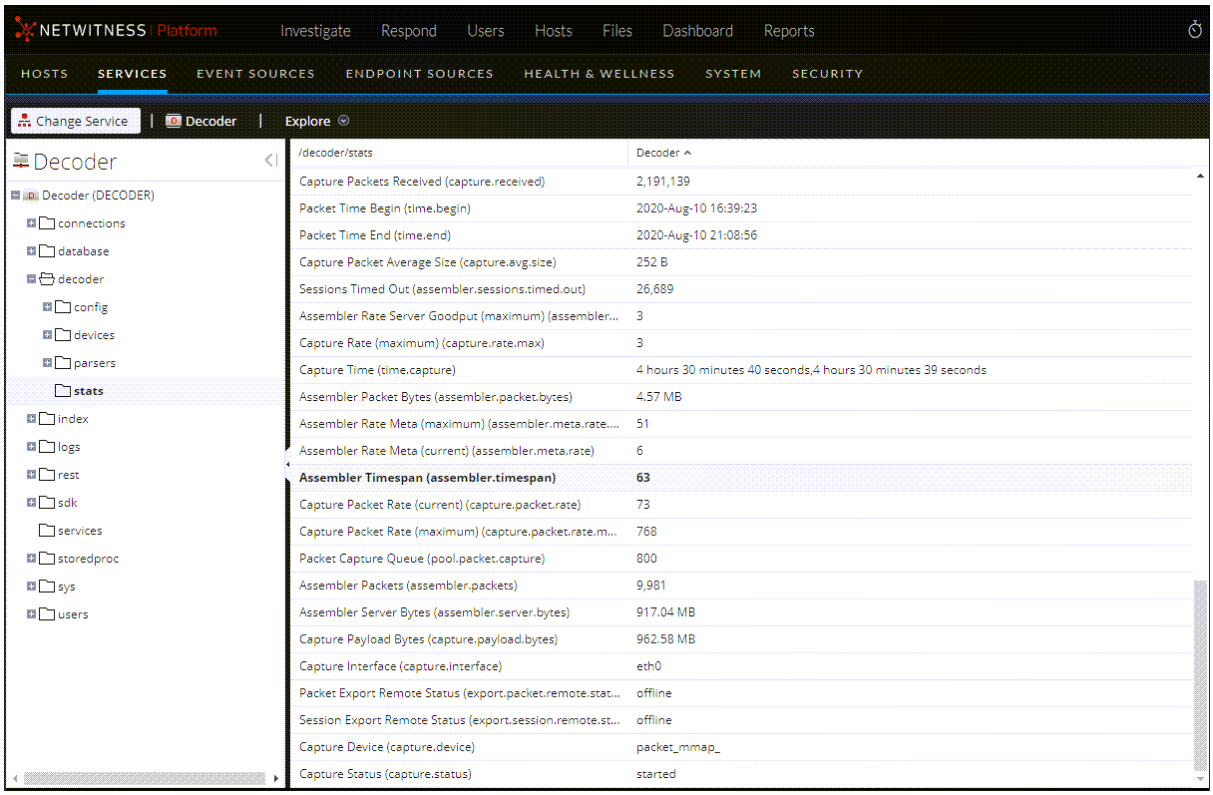
To configure Time Split mode, set the following configuration parameters and restart the Decoder:

1. In the (Admin) > **Services** view, select the Decoder service and **> View > Explore.**
2. In the Services Explore view select **decoder** > **config**.



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** >

stats.


NETWITNESS | Platform
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
HOSTS
SERVICES
EVENT SOURCES
ENDPOINT SOURCES
HEALTH & WELLNESS
SYSTEM
SECURITY
Change Service
Decoder
Explore
Decoder
Decoder (DECODER)
connections
database
decoder
config
devices
parsers
stats
index
logs
rest
sdk
services
storedproc
sys
users
/decoder/stats
Decoder
Capture Packets Received (capture.received) 2,191,139
Packet Time Begin (time.begin) 2020-Aug-10 16:39:23
Packet Time End (time.end) 2020-Aug-10 21:08:56
Capture Packet Average Size (capture.avg.size) 252 B
Sessions Timed Out (assembler.sessions.timed.out) 26,689
Assembler Rate Server Goodput (maximum) (assembler... 3
Capture Rate (maximum) (capture.rate.max) 3
Capture Time (time.capture) 4 hours 30 minutes 40 seconds,4 hours 30 minutes 39 seconds
Assembler Packet Bytes (assembler.packet.bytes) 4.57 MB
Assembler Rate Meta (maximum) (assembler.meta.rate.... 51
Assembler Rate Meta (current) (assembler.meta.rate) 6
Assembler Timespan (assembler.timespan) 63
Capture Packet Rate (current) (capture.packet.rate) 73
Capture Packet Rate (maximum) (capture.packet.rate.m... 768
Packet Capture Queue (pool.packet.capture) 800
Assembler Packets (assembler.packets) 9,981
Assembler Server Bytes (assembler.server.bytes) 917.04 MB
Capture Payload Bytes (capture.payload.bytes) 962.58 MB
Capture Interface (capture.interface) eth0
Packet Export Remote Status (export.packet.remote.stat... offline
Session Export Remote Status (export.session.remote.st... offline
Capture Device (capture.device) packet_mmap_
Capture Status (capture.status) started

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 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 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**.

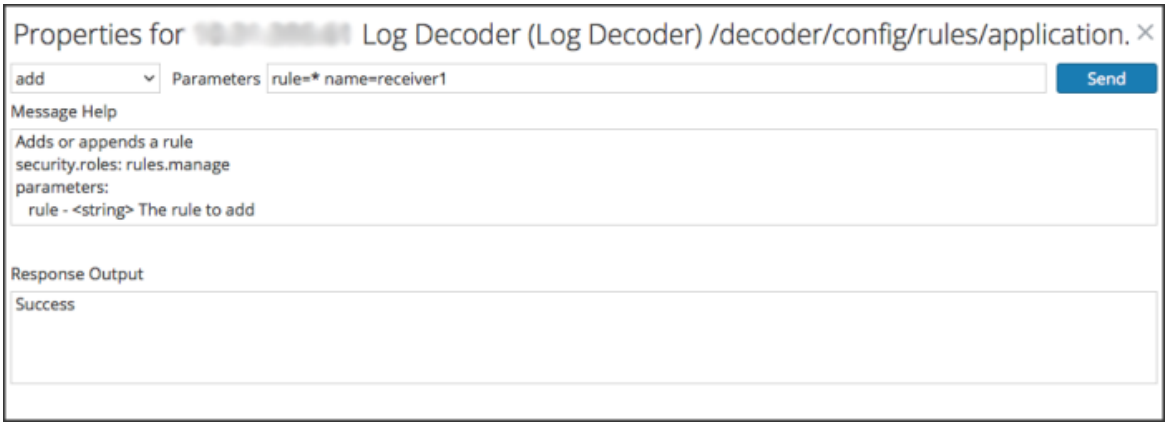
   c. In the **Properties** view, specify the **add** command with the following parameters:

   `rule=<query> name=<name>`

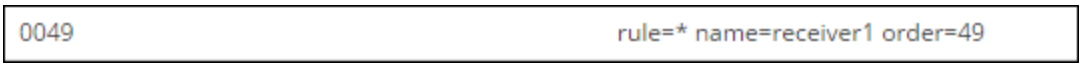
   Example 1: `rule=* name=receiver1`

   Example 2: `rule="device.type='winevent_nic'" name=receiver`)

   d. Click **Send**.



NetWitnesscreates the `name=receiver1 rule=* order=<n>` rule. NetWitness inserts the order number (for example, `order=49`) based on when you set up the rule.



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 and 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`



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 NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

c. In the `/decoder/config/logs.forwarding.enabled` parameter, specify **true**.



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:

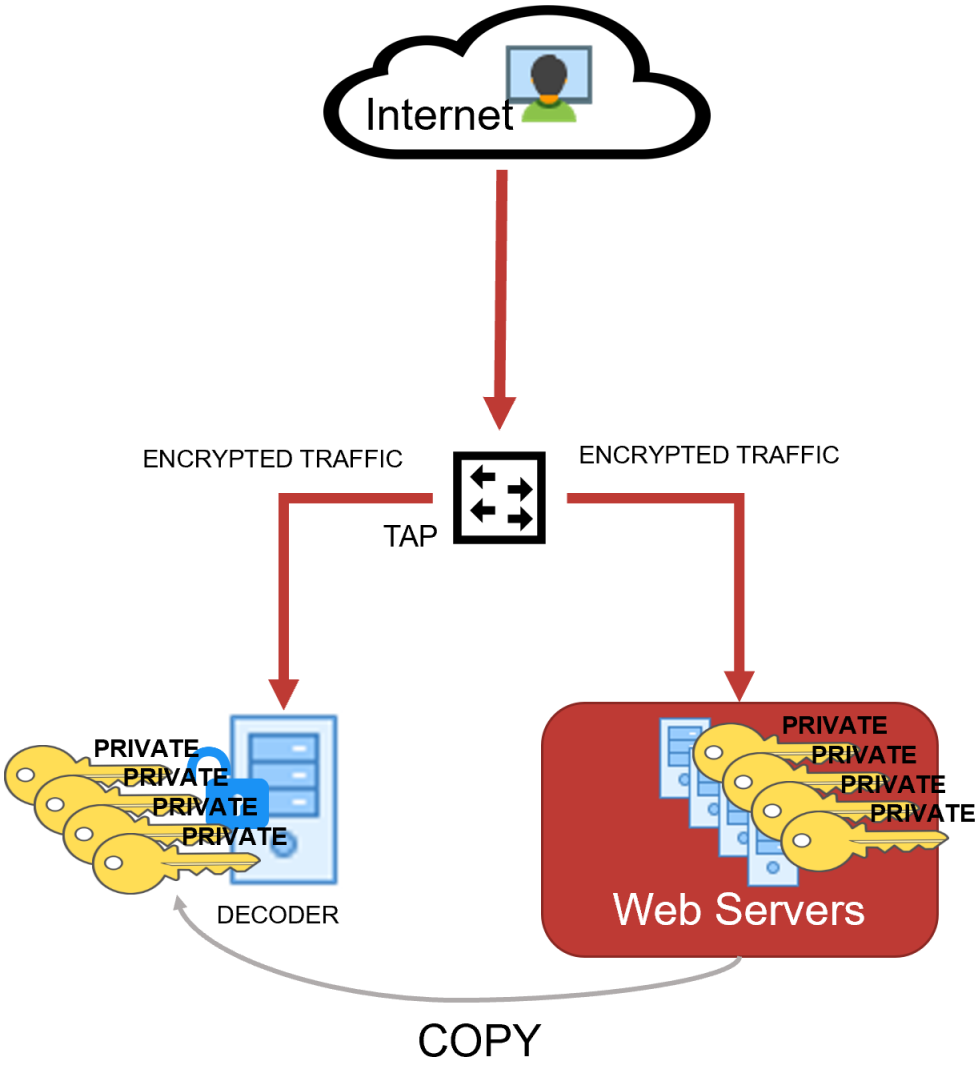
1. Go to  **(Admin) > Services** and select a Decoder or Log Decoder.
2. Select  **> View > Config**.
3. In the Configuration view, select the **Data Export** tab.
4. In the Data Export list, click the name of the host.

The Event Source tab is displayed with export connector typespec selected.

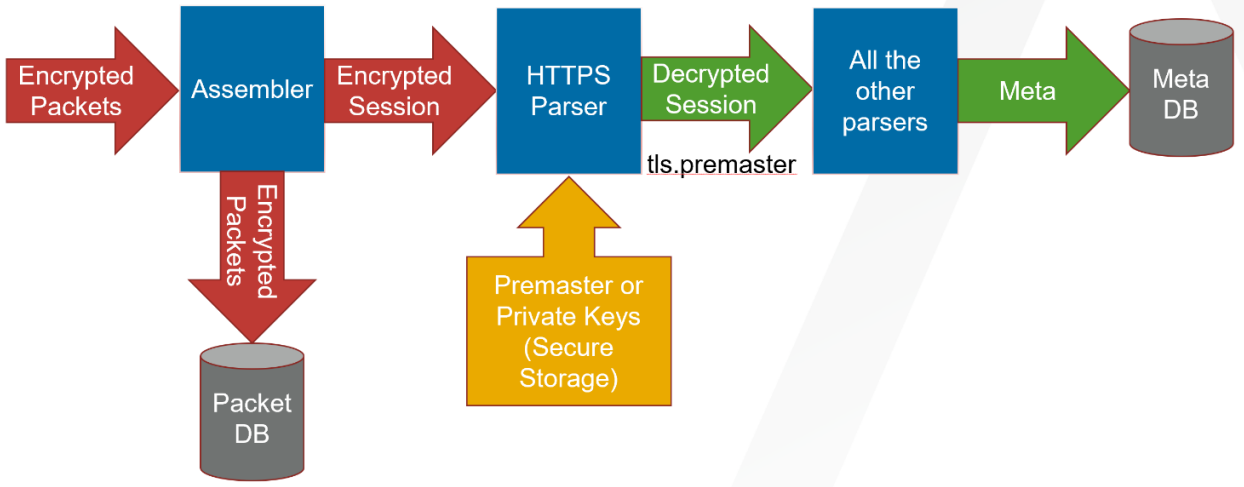
For more information on configuring export connector, see Configure Logstash Event Sources in NetWitness.

Decrypt Incoming Packets (TLS 1.2)

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 decypted 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.

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.

Extracting Certificates for the Decoder SSL Decryption

Extracting the proper key materials from popular servers is required to leverage the Decoder built-in decryption capability. The following are a few examples to extract private keys and certificates from popular servers like Apache, Microsoft IIS, and Exchange.

Note: These procedures can also be used as references to extract keys from other servers as well.

Extract private key from Apache

The private key file location would be referenced in the main Apache configuration file, which is `httpd.conf` or `apache2.conf`. The config directive `SSLCertificateKeyFile` will specify the path on your server where private keys are stored.

If Apache is using OpenSSL default certificates and keys, then locate private keys (which are usually saved to the folder /usr/local/ssl by default). You can run the command openssl version –a to find OPENSSLDIR where your server is saving the keys.

If the private key is password protected, then it can be extracted using the below command with the password.

```
openssl rsa -in yourcertificate.pem -out serverkey.pem
```

Validate the located private key using the following command.

```
openssl rsa -in serverkey.pem -check -noout
```

Extract Private Key and Certificate from Windows Servers

Windows servers provide the interface to extract certificates and private keys in PKCS12 format files (.pfx format). It is recommended to use password protection while exporting private keys.

After exporting keys to the pfx file, you can use the OpenSSL tool to extract private keys as PEM files and upload them to Decoder.

Step 2 (a): Extract PKCS12 file from Internet Information Services (IIS)

There are multiple ways to extract certificates from IIS and the following are few examples.

Using IIS Manager App

- Open Internet Information Services Manager App.
- Under IIS > Navigate to Select Certificates > Select Features view > Select the required web or application server certificate.
- Right-click the certificate and select the Export option. In the export wizard, select the file path to save the pfx file (PKCS12 format)
- Example: C:\ServerCerts\servercrt.pfx.
- Enter the password and confirm the password in the export.
- This would create password protected pfx file which contains a private key and associated certificate.

Using Microsoft Management Console (MMC)

- Open Microsoft Management Console (MMC).
- In the Console Root expand Certificates (Local Computer).
- If the Certificates folder is not available then navigate to File > Add Snap In > Add Certificates.
- Locate your server certificate, it is located in the Web Server or Personal sub-folders.
- Right-click the certificate, identified by the Common Name, select Export, and follow the wizard. Use password to protect the pfx file.
- This would create password protected pfx file which contains a private key and associated certificate.

Step 2 (b): Extract PKCS12 file from Exchange Servers

Refer the Export a Certificate from an Exchange Server document for information on how to extract the pfx file from the Exchange server.

Step 2 (c): Extract Private Key From a pfx File Using OpenSSL

OpenSSL tool provides API to export keys in PEM format.

1. Extract private key in PEM format, at the openssl prompt use the password to decrypt the pfx file.
   ```
   openssl pkcs12 -in servercrt.pfx -nocerts -nodes -out serverkey.pem
   ```
2. Validate the private key generated
   ```
   openssl rsa -in serverkey.pem -check -noout
   ```
3. Extract cert in PEM format to validate compatibility with private key
   ```
   openssl pkcs12 -in servercrt.pfx -nokeys -out serverkey.crt
   ```
4. Validate private key and certificate modulus. The modulus should be the same for both the private key and its certificate.
   ```
   openssl rsa -noout -modulus -in serverkey.pem
   openssl x509 -noout -modulus -in serverkey.crt
   ```

Recommendations

Once private keys are extracted from servers and imported to Decoder, it is recommended to secure or delete the extracted PEM file of the private key (serverkey.pem) generated by the above OpenSSL commands.

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 NetWitness Community (https://community.netwitness.com/t5/netwitness-discussions/bd-p/netwitness-discussions).

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

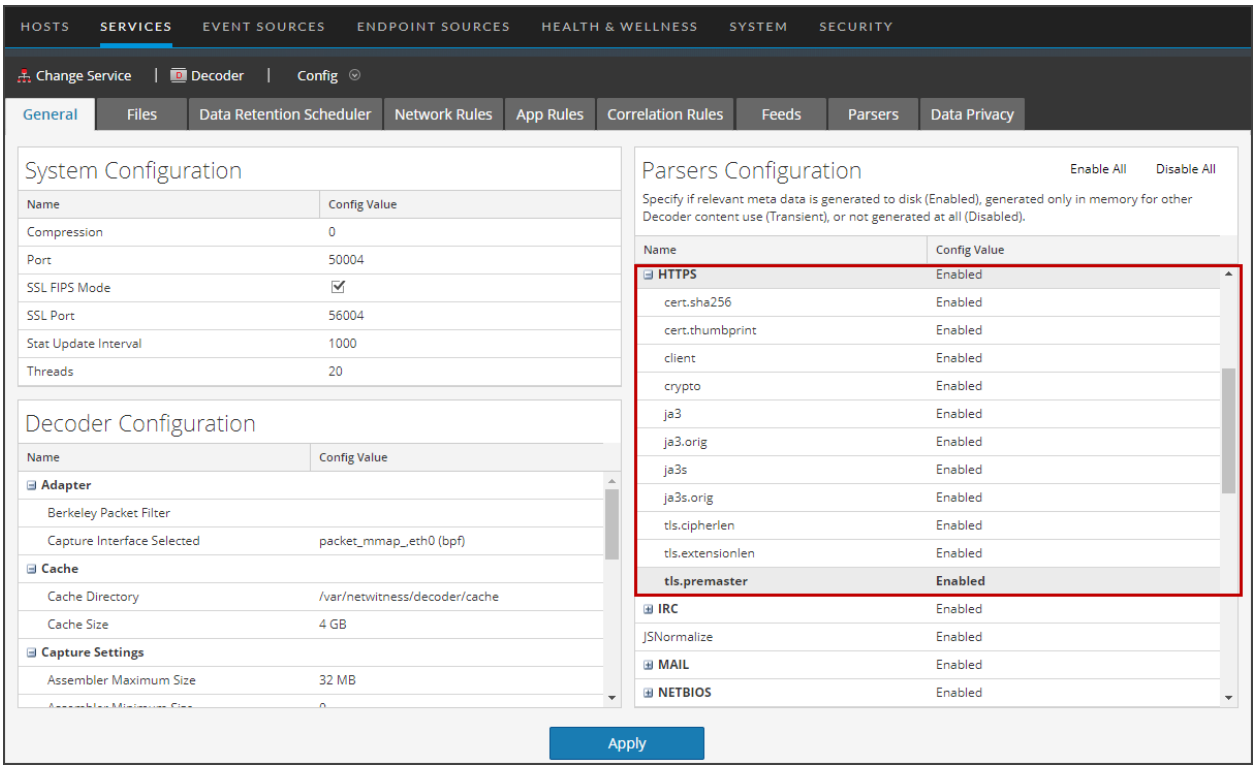
Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	ECDHE-RSA-DES-CBC3-SHA	[0xc012]	SSLv3	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	ECDHE-ECDSA-DES-CBC3-SHA	[0xc008]	SSLv3	Kx=ECDH	Compliant
TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	EDH-RSA-DES-CBC3-SHA	[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

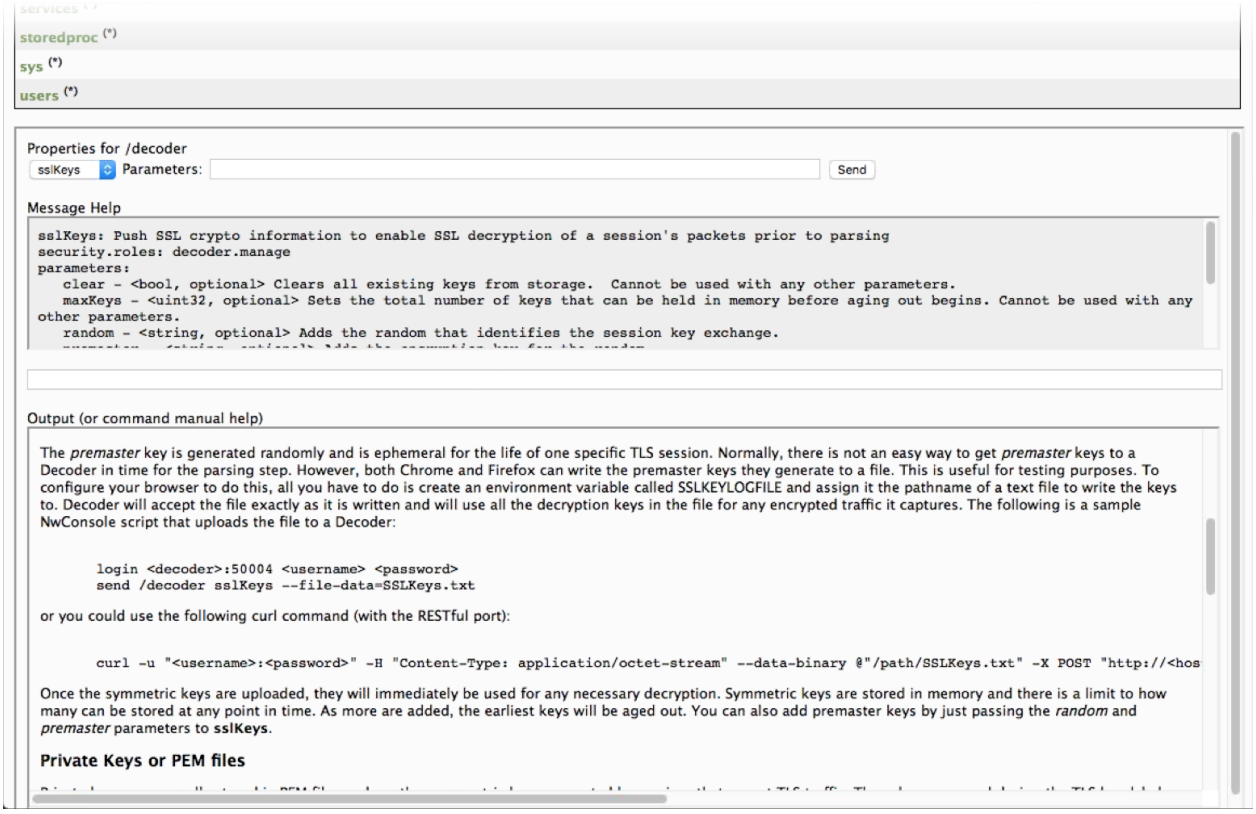
The decryption process requires the native HTTPS parser to be enabled to examine the TLS handshake and to compare the uploaded encryption keys to determine if they decrypt the traffic. This can be validated in (Admin) > **Services** > **Decoder** > **Config** > Parsers Configuration panel as shown in the following figure.



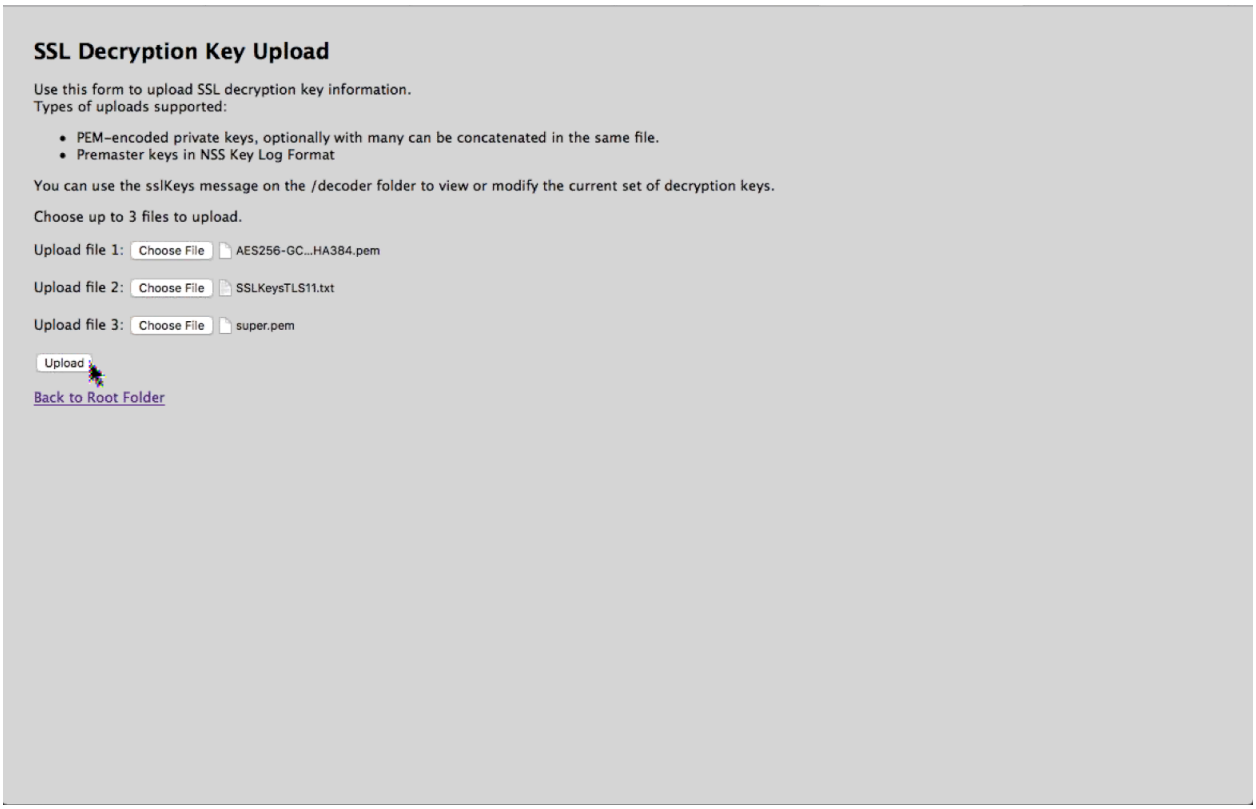
Step 5: Upload the Supported Private Keys to Decoders

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 NetWitness Platform*. For information about how to use the RESTful interface, see the *RESTful API User Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.



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.

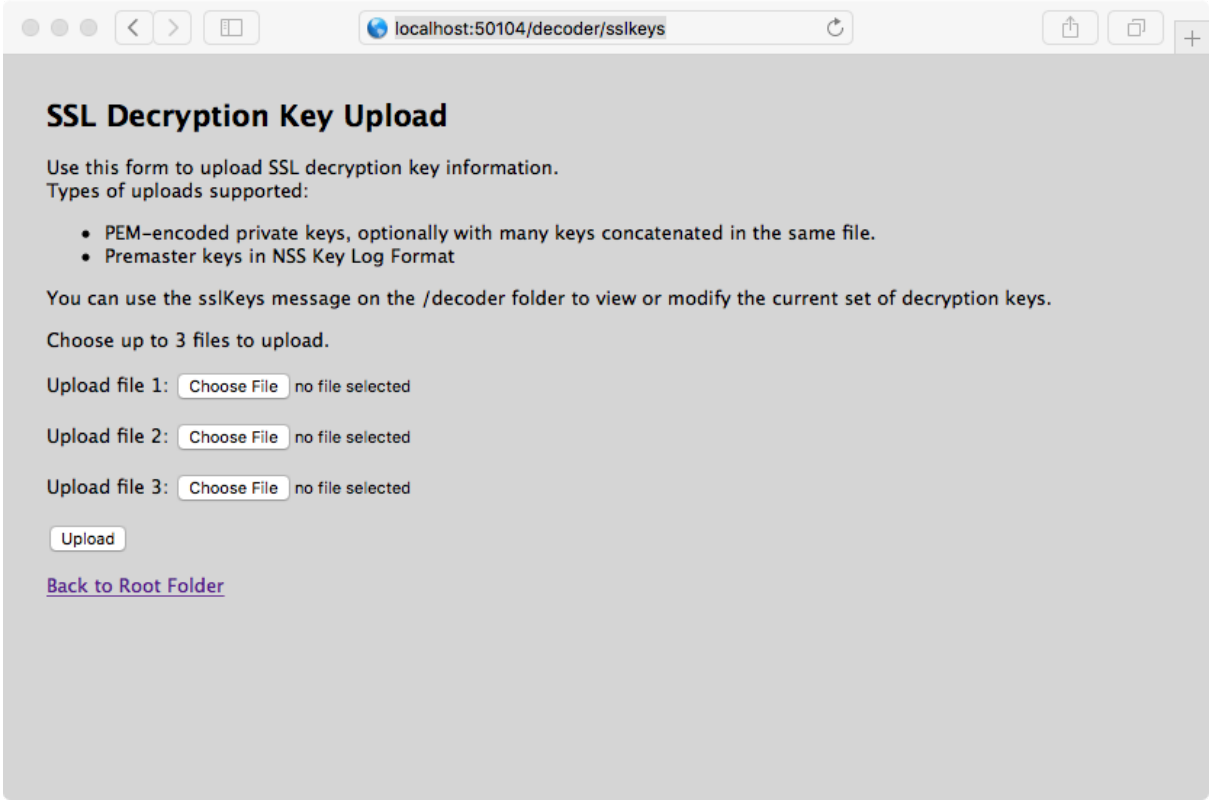


Upload Multiple Premaster and Private Keys

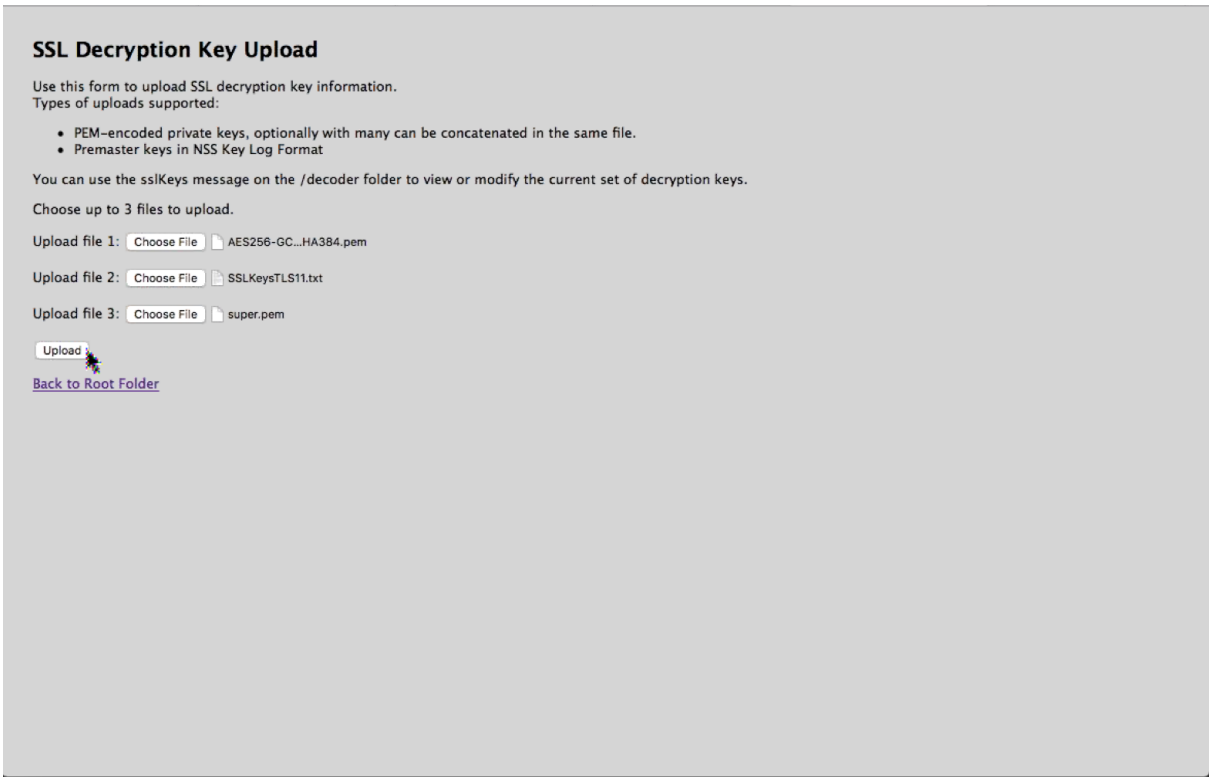
Note: As of current version, uploading premaster keys is not officially supported yet for real-time capture, but can be used to validate that decryption is functioning properly on the Decoder.

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`.

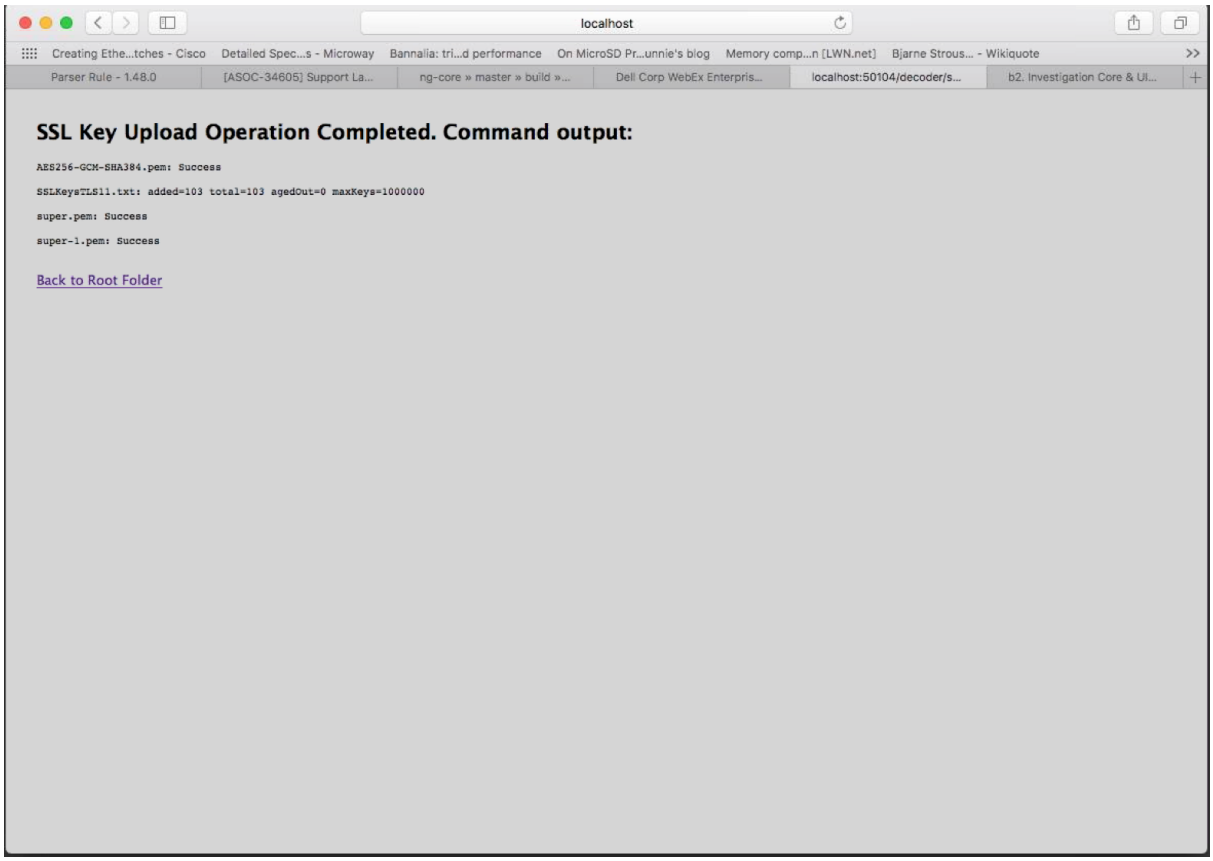


2. Next to **Upload File 1**, click **Choose File** and locate the premaster key file or PEM file that you want to upload on your local file system.
3. (Optional) Repeat for **Upload File 2** and **Upload File 3**.



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
`clear`	Removes all premaster keys from memory. Does not delete any PEM files installed on the system.
`maxKeys`	Changes the maximum number of premaster keys that are stored in memory.
`listPems`	Returns a list of all installed private key PEM files.
`deletePem`	Deletes the named PEM file from the file system. You can pass this parameter more than once to remove multiple files.
`random`	The random hash used to identify the premaster key.
`premaster`	The premaster key that will be installed for the previous `random` parameter. They must show up in pairs and `random` 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"
```

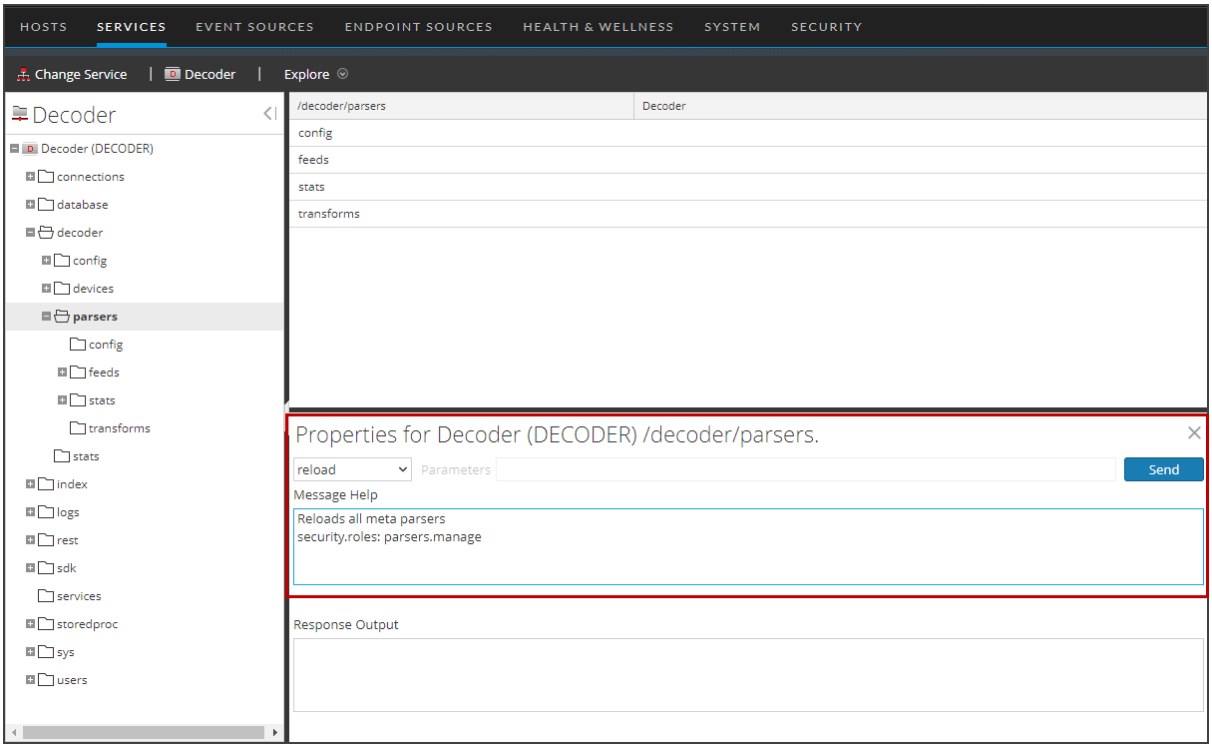
Step 6: Reload HTTPS Parser on Decoders

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 formated 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

NetWitness 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.

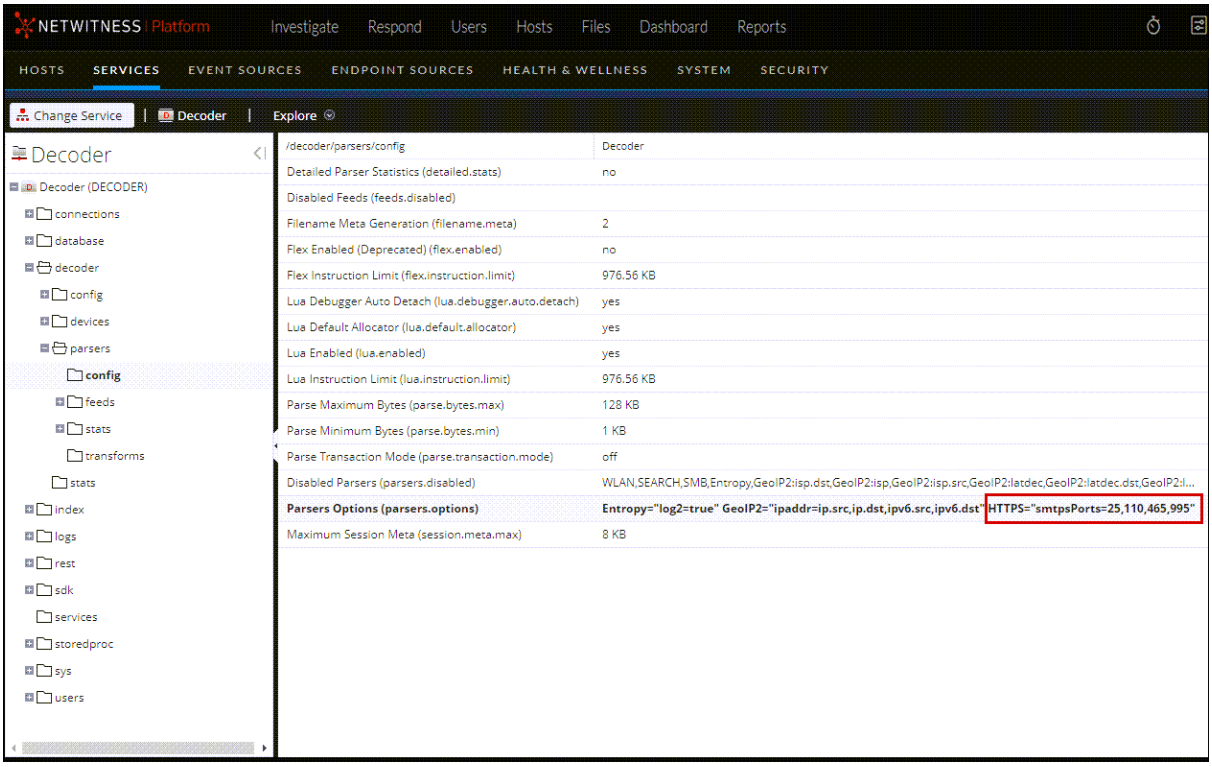
The STARTTLS search is limited to the first 1024 bytes of the session's payload, to prevent performance degradation. If STARTTLS is found, the parser uses the next client packet with payload as the start of a TLS negotiation.

To add an HTTPS parser option for secure SMTP decryption:

1. In the NetWitness 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.



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 Assembler 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)

Parser Name	Description
TNS	Oracle database protocol (TNS)
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.sha1=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.

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`.

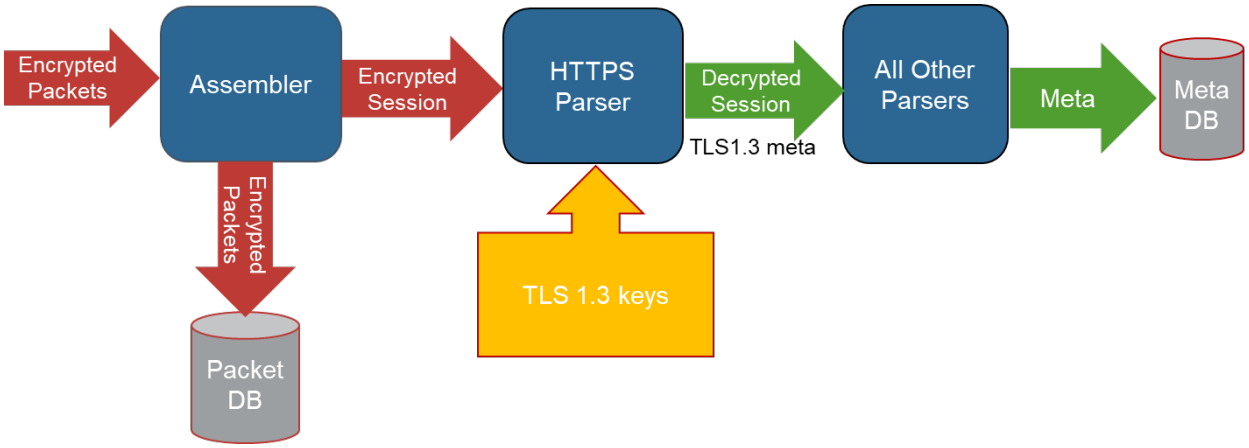
The Community ID is written to a text-formatted meta key named `community`.

The Community ID is generated on sessions containing TCP over IPv4, TCP over IPv6, UDP over IPV4, and UDP over IPV6.

Decrypt Incoming Packets TLS 1.3

Administrators can configure a Network Decoder to decrypt incoming TLS 1.3 network packets. This capability is currently supported with encrypted communications to managed servers using TLS 1.3 or earlier.

For TLS 1.2 and earlier protocols refer Decrypt Incoming Packets (TLS 1.2)



As the network traffic is collected and passed through the system, the encrypted packets are stored as captured even if the system can decrypt them. As the Decoder parses the encrypted traffic, it uses TLS 1.3 keys that have been pushed or uploaded, and determines if the network traffic can be decrypted. If the Decoder is able to decrypt the traffic, it generates an additional meta keys, defined as `tls.client.hdshk, tls.server.hdshk, tls.client.app, tls.server.app, tls.client.early` which contains the decryption keys 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.

TLS 1.3 Keys and Cipher suites

According to RFC 8446 the TLS 1.3 protocol doesn't support Private key based decryption and it only support a set of ephemeral symmetric keys also called as key secrets for the life of one specific TLS 1.3 session. The TLS 1.3 session also encrypts handshake, all the TLS records after Server Hello are encrypted , and at each stage a corresponding key is used for the encryption ex: client handshake, server handshake, client application and server application etc.. . So in order for decrypting a TLS 1.3 session all the required session keys `tls.client.hdshk, tls.server.hdshk, tls.client.app, tls.server.app, tls.client.early`(optional) are needed to be available.

Cipher Suites

TLS 1.3 protocol recommends to support only strong and advanced Cipher suites. The Network Decoder supports all the advertised cipher suites from RFC-8446, however on integration with Thirdparty key forwarder the supported cipher suites might be limited to list supported by Thirdparty server. Example only GCM ciphers are supported etc..

The following is list of cipher suited supported in Network Decoder

- TLS_AES_128_GCM_SHA256
- TLS_AES_256_GCM_SHA384
- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_128_CCM_SHA256
- TLS_AES_128_CCM_8_SHA256

Steps to configure Network Decoder to decrypt TLS 1.3 traffic

- Step 1: Validate that the Network Decoder Captures Encrypted TLS 1.3 Traffic
- Step 2: Confirm HTTPS Parser is Enabled on Decoders
- Step 3: Integrate Network Decoder to receive keys from Thirdparty TLS 1.3 key forwarder
- Step 4: (Optional) Upload TLS 1.3 keys to validate Decryption
- Step 5: Validate TLS 1.3 keys are received on Decoder
- Step 6: View Decrypted Sessions

Step 1: Validate that the Network Decoder Captures Encrypted TLS 1.3 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
```

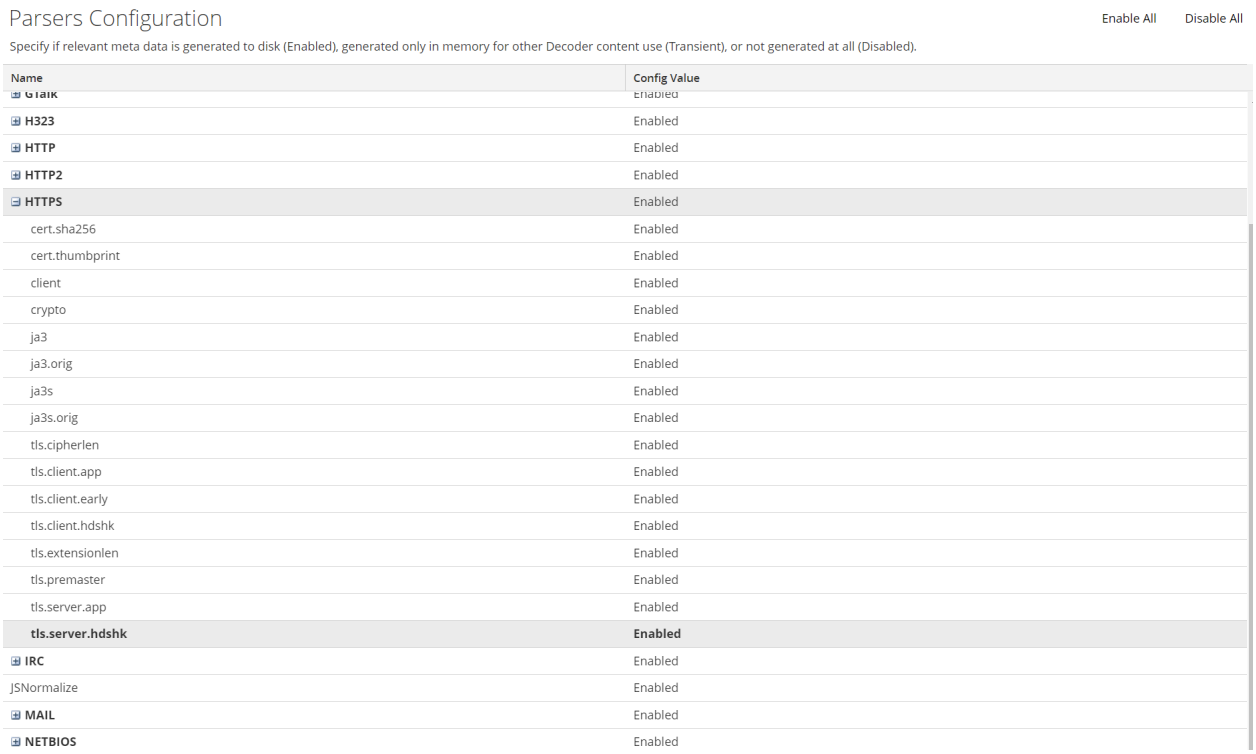
Conduct this search on the Concentrator that aggregates metadata from the Decoder that is expected to capture this network segment, or on a Broker that may have access to all collected metadata.

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: 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.



If TLS_Lua parser is enabled then meta version: TLS 1.3 is registered for TLS 1.3 encrypted traffic.

Step 3: Integrate Network Decoder to receive keys from Thirdparty TLS 1.3 key forwarder

TLS 1.3 protocol only support ephemeral keys which are unique per session and Network Decoder need to receive those keys in time for parsing step.

Network Decoder provides `sslKeys` API which can receive the keys forwarded or injected to Decoder on RESTful interface.

A Thirdparty TLS 1.3 key forwarder which acts as Man-In-The-Middle (MITM) gateway for TLS 1.3 may intercept SSL handshake and forward the TLS 1.3 keys to Network Decoder.

Example: F5 BIG-IP Local Traffic Manager.

Option 1: Configure F5 BIG-IP Local Traffic Manager (LTM)

If you are using BIG-IP LTM in your environment for application traffic management, load balancing etc.. and if you would like to decrypt and analyze TLS 1.3 traffic handled by BIG-IP LTM, then refer the following section to Configure F5 BIG-IP LTM and forward keys to Network Decoder.

For more information refer: F5 BIG IP - Netwitness Perfect Forward Secrecy Inspection Visibility

Option 2: Using sslKeys API

If you are using other Man-In-The-Middle (MITM) gateway which can intercept TLS 1.3 handshake and forward the keys, then you can refer the following information and parameters to send keys to `sslKeys` API on Network Decoder.

The `sslKeys` API is used to push SSL crypto information to enable SSL decryption of a session's packets prior to parsing.

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
clear	Removes all premaster and TLS 1.3 keys from memory. Will not delete any PEM files installed on the system.
maxKeys	Changes the maximum number of premaster and TLS 1.3 keys that will be stored in memory.
counts	Returns key counts for sslKeys. Cannot be used with any other parameters.
listPems	Returns a list of all installed private key PEM files.
deletePem	Deletes the named PEM file from the file system. You can pass this parameter more than once to remove multiple files.
random	The random hash used to identify the premaster key or TLS 1.3 keys.
premaster	The premaster key that will be installed for the previous random parameter. They must show up in pairs and random must be first.
CETS	The TLS 1.3 client early traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.
CHTS	The TLS 1.3 client handshake traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.
SHTS	The TLS 1.3 server handshake traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.
CTS0	The TLS 1.3 client traffic secret 0 key (application) that will be installed for the previous random parameter. It must show up with random parameter.
STS0	The TLS 1.3 server traffic secret 0 key (application) that will be installed for the previous random parameter. It must show up with random parameter.

Usage:

The Client random and TLS 1.3 keys can be passed as parameters to `sslKeys`, ex:

- Using Native port 56004:
  - `/decoder sslKeys random=`<value> `CETS=`<value> `CHTS=`<value> `SHTS=`<value> `CTS0=`<value> `STS0=`<value>
- Using RESTful port 50104:
  - `/decoder?msg=sslKeys&random=`<value>`&CETS=`<value>`&CHTS=`<value>`&SHTS=`<value>`&CTS0=`<value>`&STS0=`<value>
- Replace <value> with base64 encoded TLS 1.3 key value in the above cases.

Note: For most of the sessions the CETS (CLIENT_EARLY_TRAFFIC_SECRET) key may not be generated or used by TLS clients and it is normal. However, if the other keys (handshake and application keys) are not generated and not available then TLS 1.3 decryption will NOT be successful.

Return Values

Most commands return name/value pairs of statistics about the premaster and TLS 1.3 keys in memory. The statistics are listed in the following table.

Parameter	Description
added	The number of premaster and TLS 1.3 keys just added during this command.
total	The total number of premaster and TLS 1.3 keys loaded in memory.
agedOut	The total number of premaster and TLS 1.3 keys that were removed during this command (this is not a lifetime stat).
maxKeys	The current maximum allowed premaster and TLS 1.3 keys.
premaster	The total number of premaster keys loaded in memory.
tls13	The total number of TLS 1.3 keys loaded in memory.

Step 4: (Optional) Upload TLS 1.3 keys to validate Decryption

Note: Uploading TLS 1.3 keys is not supported for real-time capture, but can be used to validate that decryption is functioning properly on the Decoder.

The `sslKeys` command provides a way to upload TLS 1.3 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. Now, all TLS handshakes that use the TLS 1.3 keys and 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 NetWitness Platform*. For information about how to use the RESTful interface, see the *RESTful API User Guide for NetWitness Platform*. Go to the *Master Table of Contents* to find all NetWitness Platform documents.

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`.

2. Click **Choose File** Next to **Upload File 1**, and locate the TLS 1.3 keys file that you want to upload on your local file system.

3. Click **Upload**. The files are uploaded to the Decoder and results are displayed in the form.

TLS 1.3 Keys

Similar to `premaster` key, capturing TLS 1.3 keys and uploading to Decoder in time for parsing is not an easy way. However, well known browsers like Chrome , Firefox and IE can write the TLS 1.3 keys they generate to a file and this is useful for testing purposes. To configure your browser to do this, create an environment variable called SSLKEYLOGFILE and set with a pathname to text file to write the keys. Decoder will accept the file exactly as it is written and will use all the decryption keys in the file for any encrypted traffic it captures.

The following is a sample NwConsole script that uploads the file to a Decoder:

```
login <decoder>:50004 <username> <password>
send /decoder sslKeys --file-data=SSLKeys.txt
```

Optionally the parameters can be passed in as list in the upload file, ex:

```
send /decoder sslKeys --file-data=SSLKeys.txt --file-format=params-list
```

or you could use the following curl command (with the RESTful port):

```
curl -u "<username>:<password>" -H "Content-Type: application/octet-stream" --
data-binary @"/path/SSLKeys.txt" -X POST
"http://<hostname>:50104/decoder?msg=sslKeys"
```

Once 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 are added, the earliest keys will be aged out. You can also add available TLS 1.3 keys by just passing the `random`, `CETS` , `CHTS` , `SHTS` , `CTS0` and `STS0` parameters to `sslKeys`.

Note: For most of the sessions the CLIENT_EARLY_TRAFFIC_SECRET key may not be generated or used by TLS clients and it is normal. However, if handshake and application keys are not generated and not available then TLS 1.3 decryption wouldn't be successful.

Step 5: Validate TLS 1.3 keys are received on Decoder

It is advised to validate if Decoder is receiving keys to Decrypt TLS 1.3 sessions. As Decryption wouldn’t be successful if the keys are not available.

The following are few ways:

1. Decoder generates the following TLS 1.3 meta per session

`tls.client.hdshk`, `tls.server.hdshk`, `tls.client.app`, `tls.server.app`, and `tls.client.early`

2. The `sslKeys` API provides an option to display current key stats using counts=true .

Example: `/decoder sslKeys counts=true`

```
premasterKeys: 478
tls13Keys: 341
total: 819
maxKeys: 1000000
```

Step 6: 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, the `tls.client.hdshk`, `tls.server.hdshk`, `tls.client.app`, `tls.server.app`, and `tls.client.early` metadata is associated with the session that has the TLS 1.3 keys required for decrypting that session. The session will also have crypto metadata denoting the cipher suite.

Viewing Unencrypted Traffic

If packets are decrypted during the parse stage, encrypted packets are written to disk, and the matching TLS 1.3 keys used for decrypting is written to the TLS 1.3 meta, analysts can view the unencrypted packets using the TLS 1.3 meta keys.

Option 1 : Using Event Analysis UI

Starting 12.0 the Event Analysis View support Displaying Decrypted and Encrypted Payloads of TLS sessions when TLS keys meta are registered.

Option 2 : Using RESTful API on Decoder service

The 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`.

If you are not able to view the encrypted sessions, check Troubleshooting. However, if you are able to see the TLS 1.3 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 if Decoder is receiving TLS 1.3 keys. Refer #Step 5 listed above.
2. Validate that your TLS 1.3 session key is using an acceptable cipher-suite.
3. Make sure your Network Decoder is capturing the traffic matching the server for which you have the TLS 1.3 keys. To check, you can upload the TLS 1.3 keys and PCAP into Wireshark and see if it is able to decrypt it.

Additional Features

The following features which are supported in TLS 1.2 are also supported in TLS 1.3 decryption.

1. TLS Certificate Hashing
2. JA3 and JA3S TLS Fingerprints
3. Performance Considerations

For more details refer Decrypt Incoming Packets (TLS 1.2)

Edit Decoder System Configuration

When a service is first added to NetWitness, 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 setting unless an NetWitness Customer Support technician advises you to change them.

System Configuration

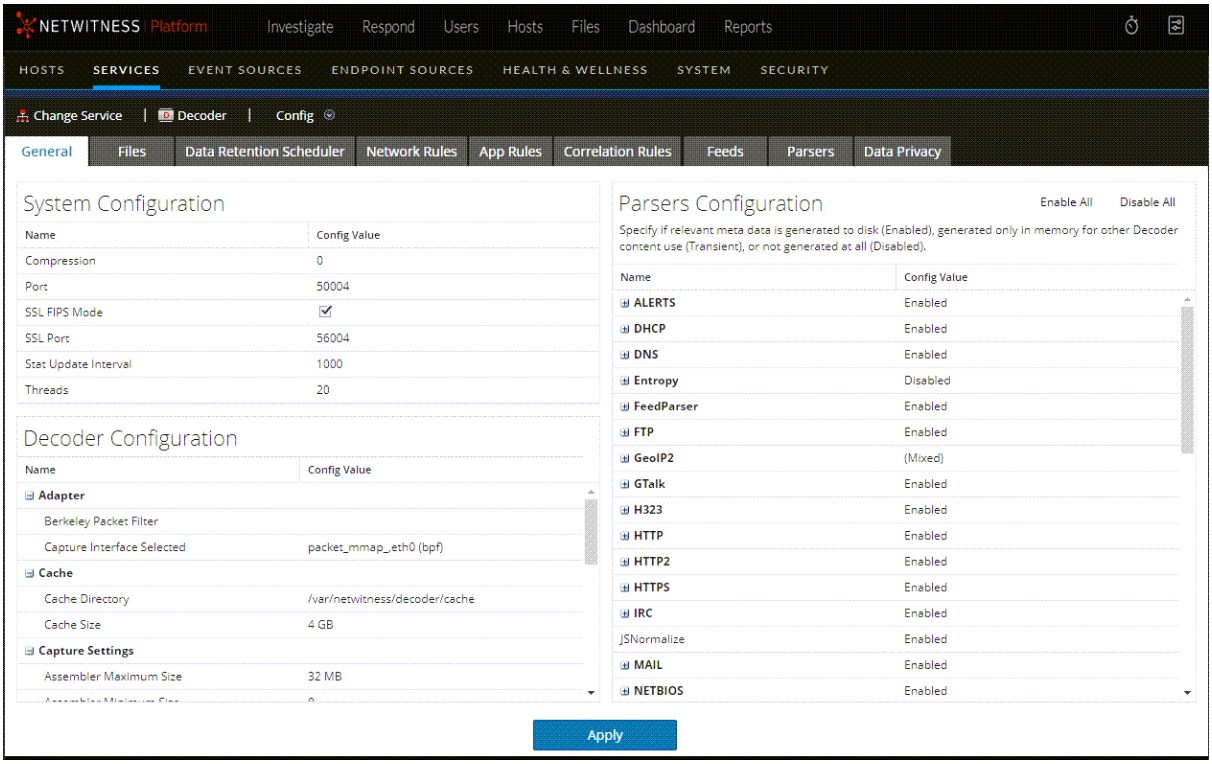
Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☐
SSL Port	56004
Stat Update Interval	1000
Threads	20

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.



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.

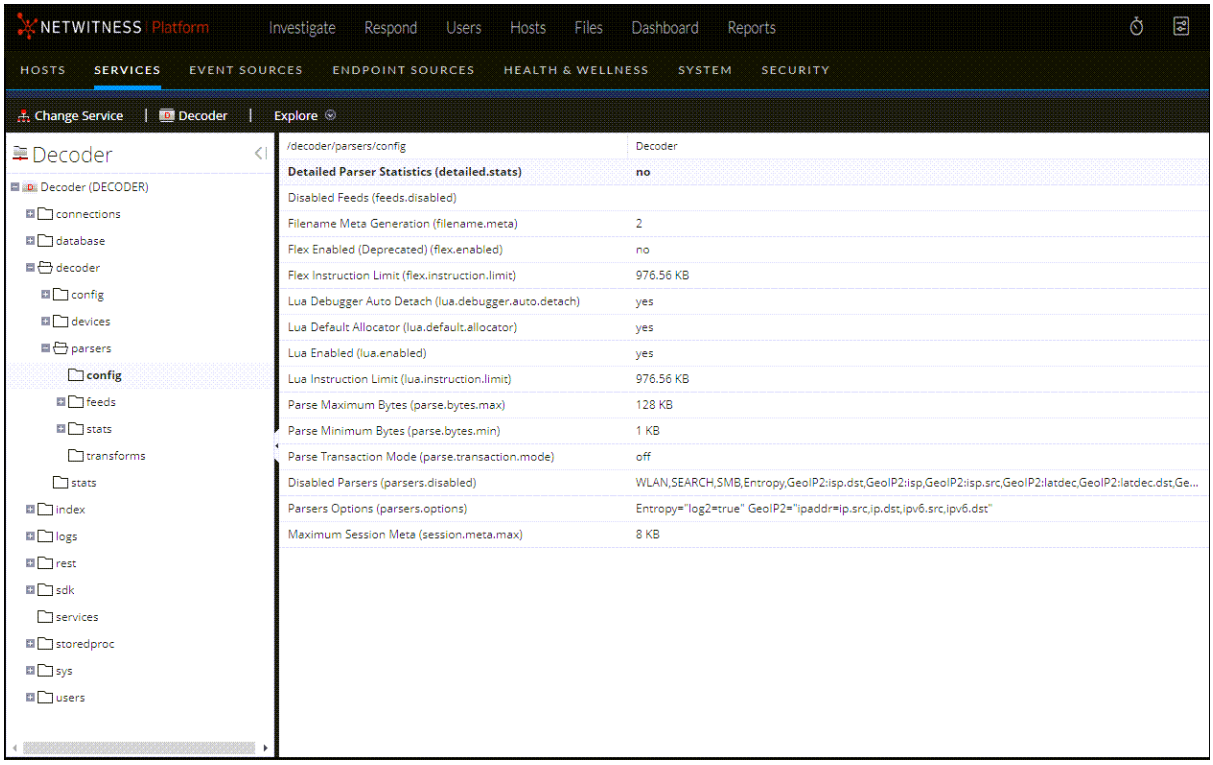
Enable CPU Usage Statistics for Installed Content

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 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.



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.

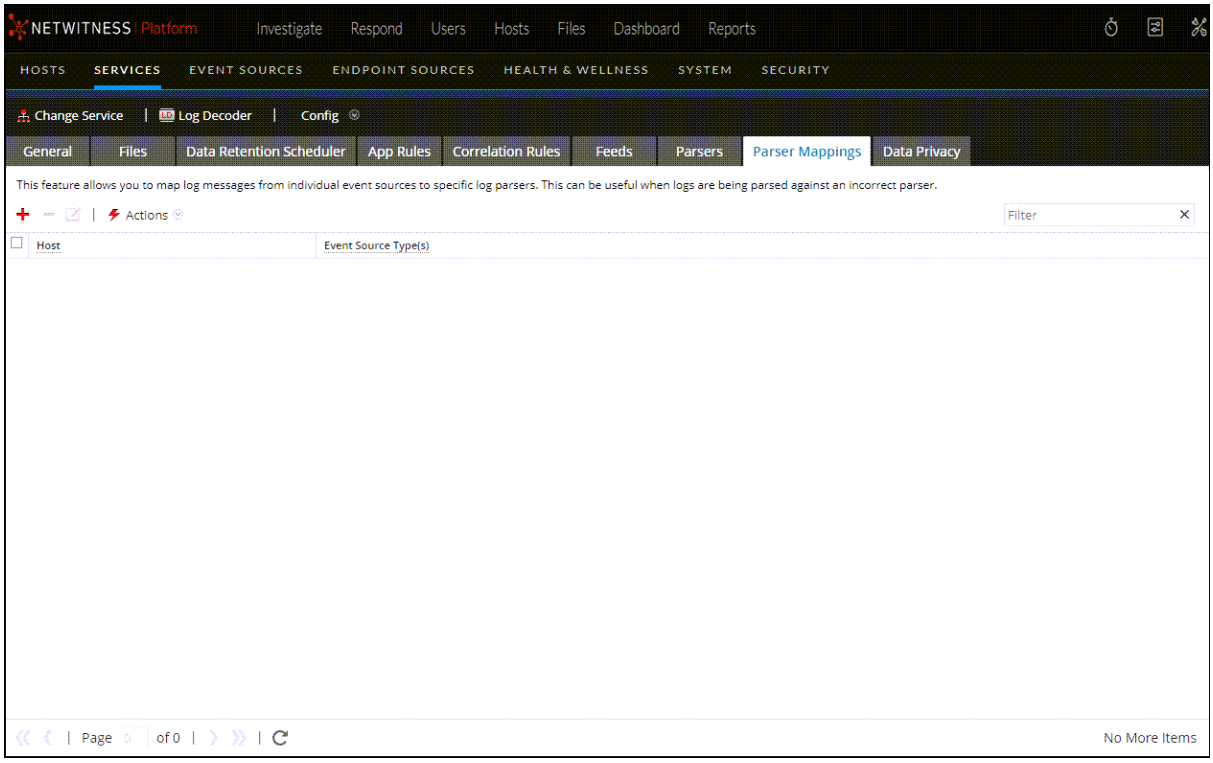
Note: You can also enable parser mapping functions by navigating to  (Admin) **> Event Sources > Discovery.**

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.



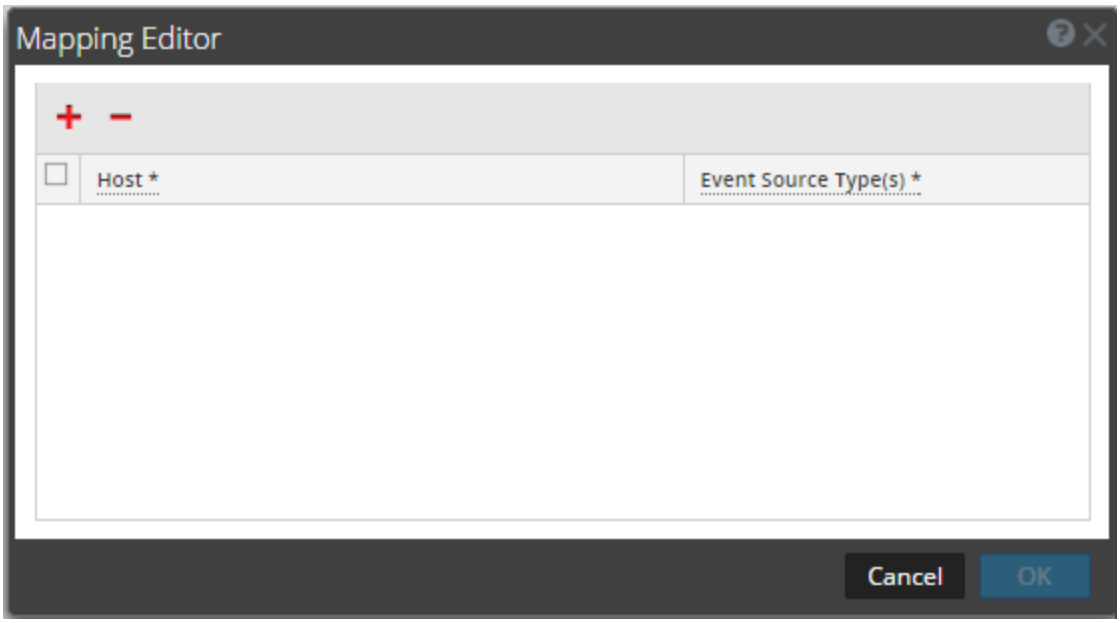
Update IP to Event Source Mapping

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 Sources(**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

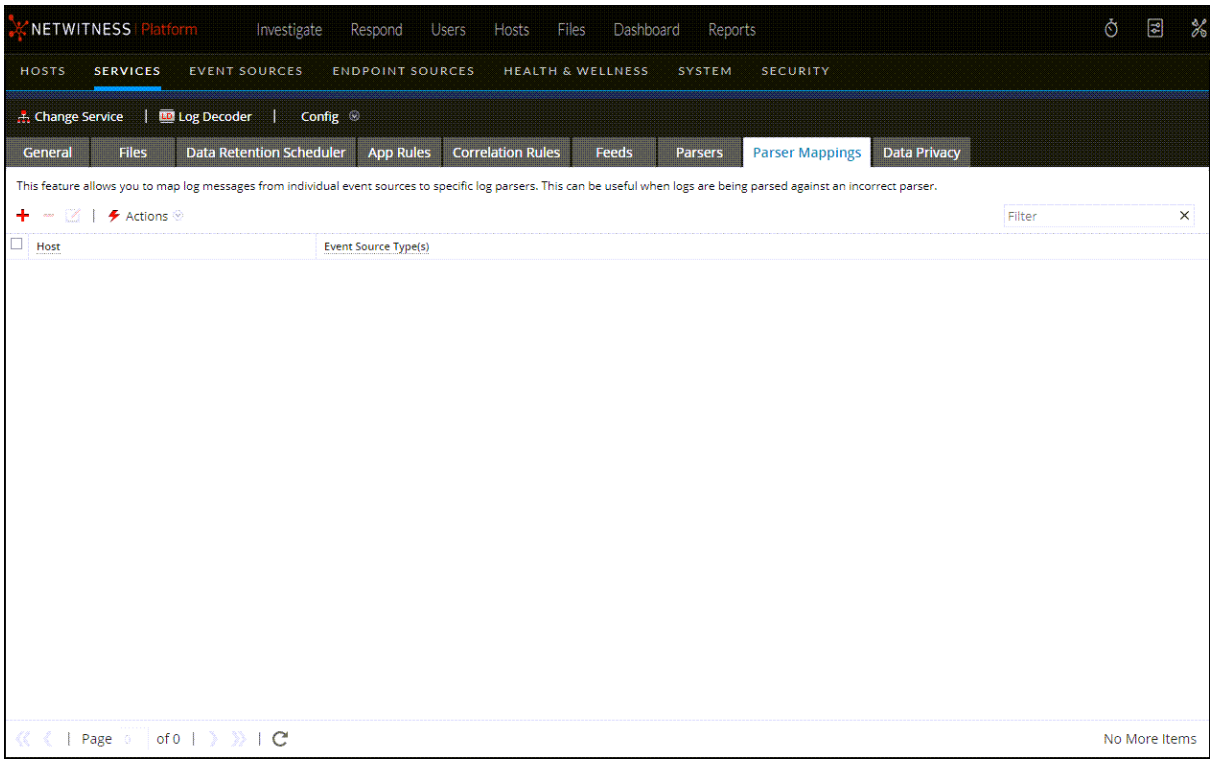
To read an IP to event source type mappings:

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.

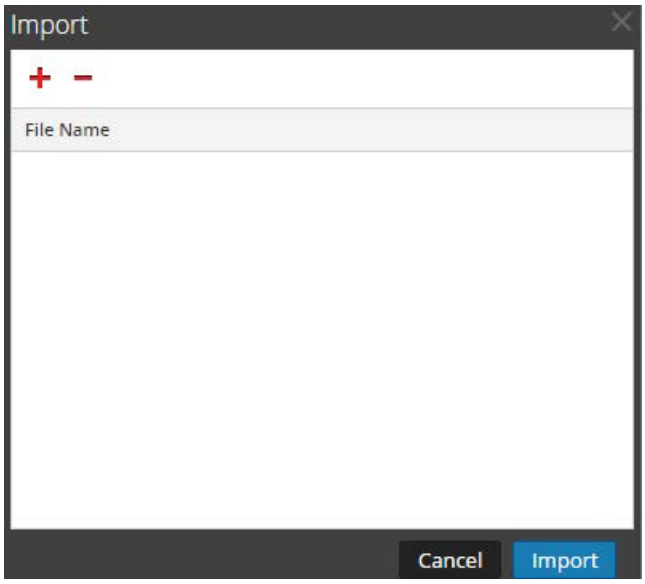
Import IP to Event Source Mapping Entries

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.



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.

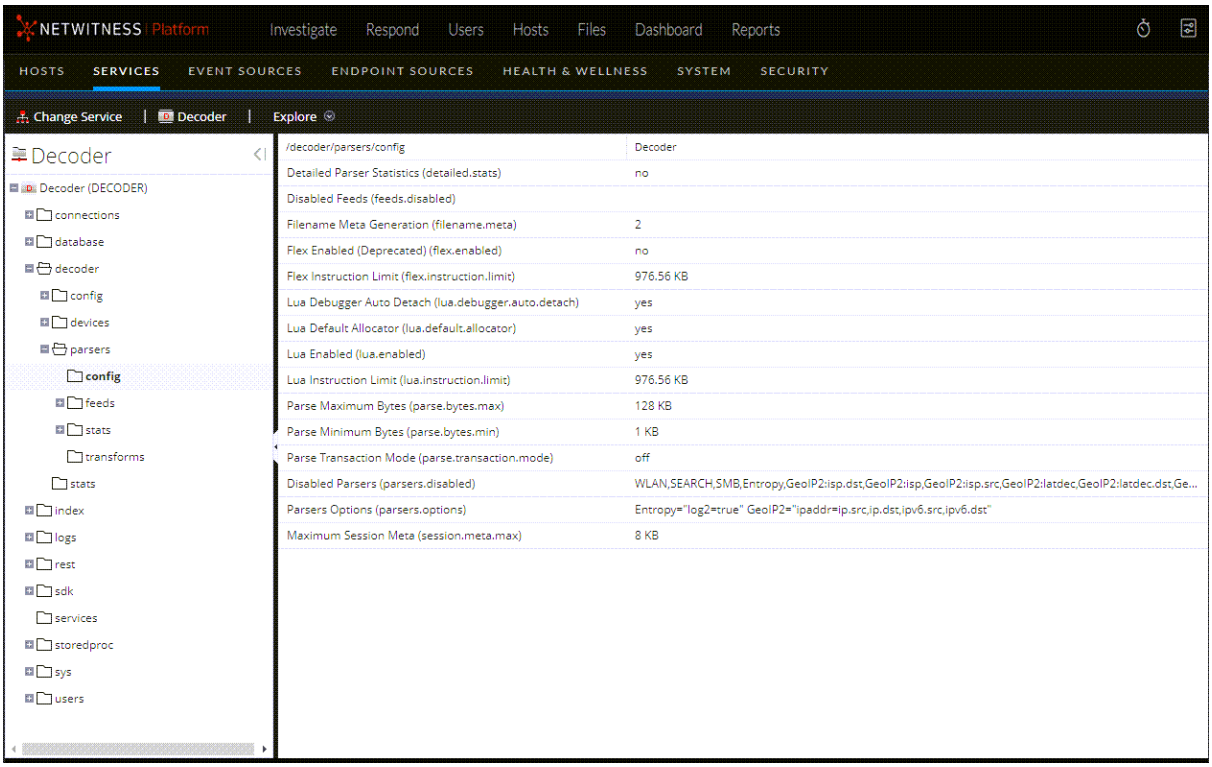
The settings to enable or disable Lua and Flex parsing systems are configured correctly by default and you do not typically have to change them. However, you may need to adjust these settings at the request of NetWitness Customer Care or for troubleshooting purposes.

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`**.



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.

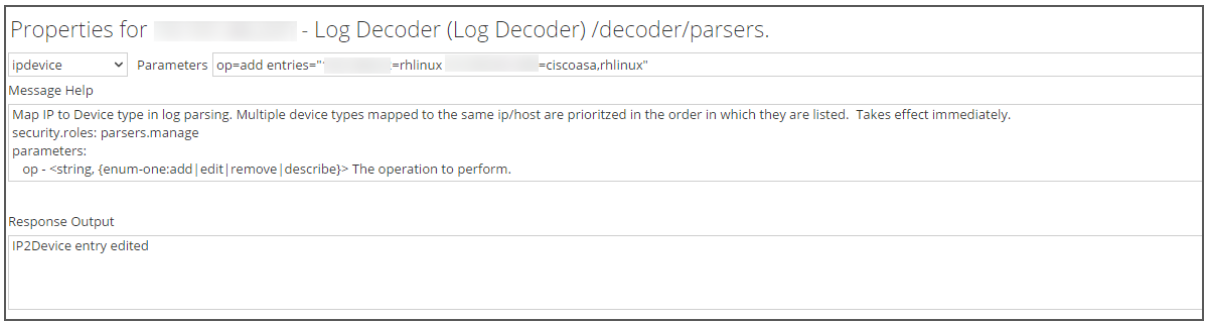
Map an IP Address to a Service Type

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**.



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 NetWitness 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 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:

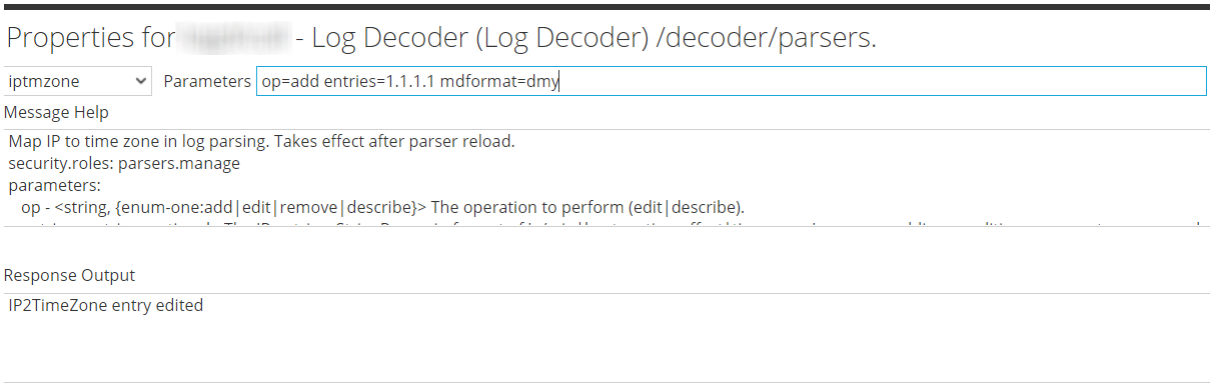
1. Go to **(Admin) > Services**.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select; **> View > Explore**.
3. Go to **/decoder/parsers** node, right click **Parsers**, and select **Properties**.
4. In the **Properties** view, select **iptmzone** from the drop-down 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**.

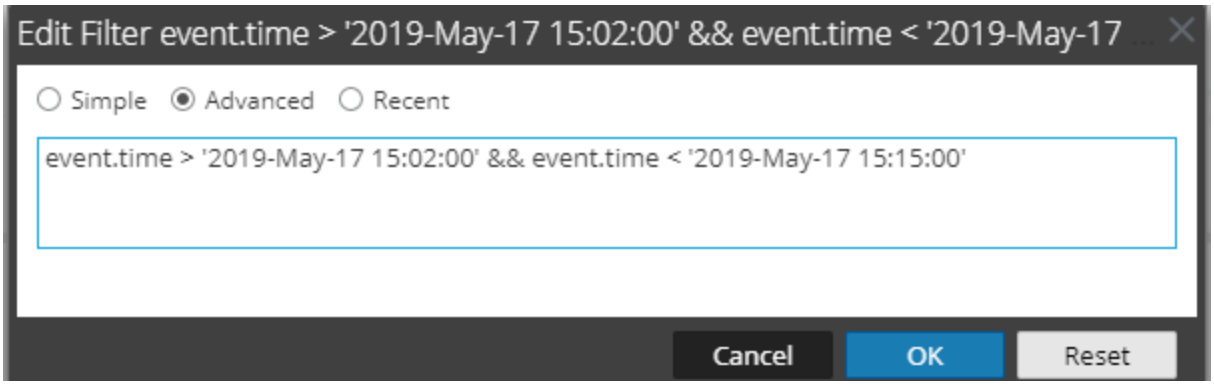


5. Click **Send**.

NetWitness maps the IP address along with the date format in the Log Decoder. Event time meta items are updated according to their respective mappings.

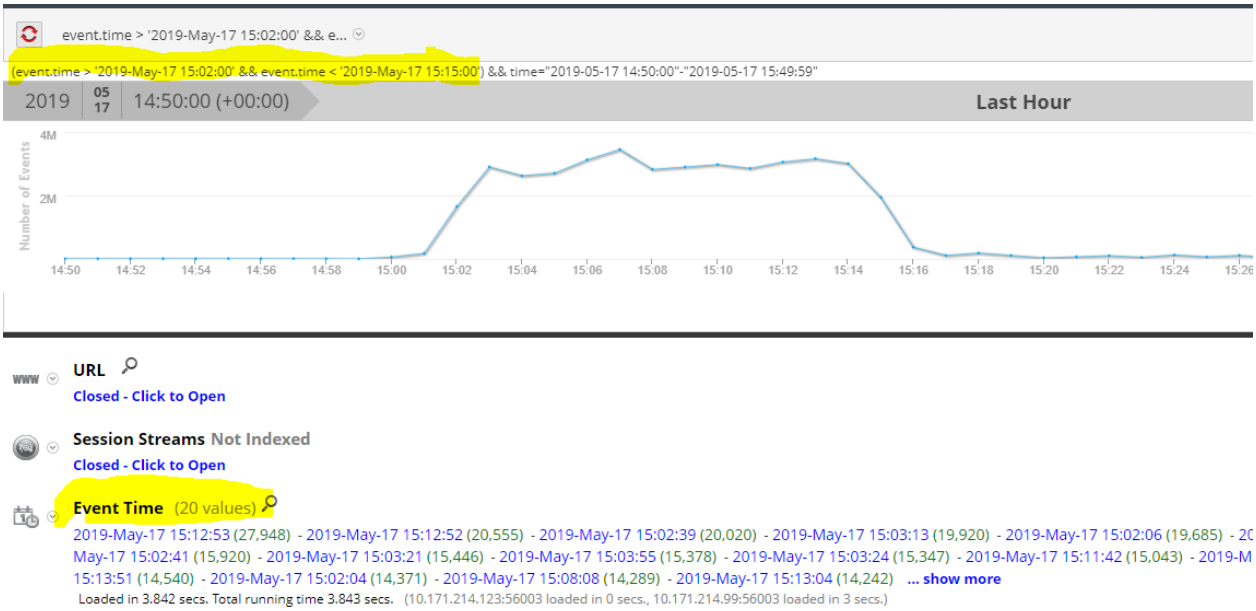
Event Time Filter Example

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:



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 Log Decoder

NetWitness added the capability to view a small sampling of recent logs for specific devices through detail tabs of the Discovery View.

To enable logs preview for an 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.

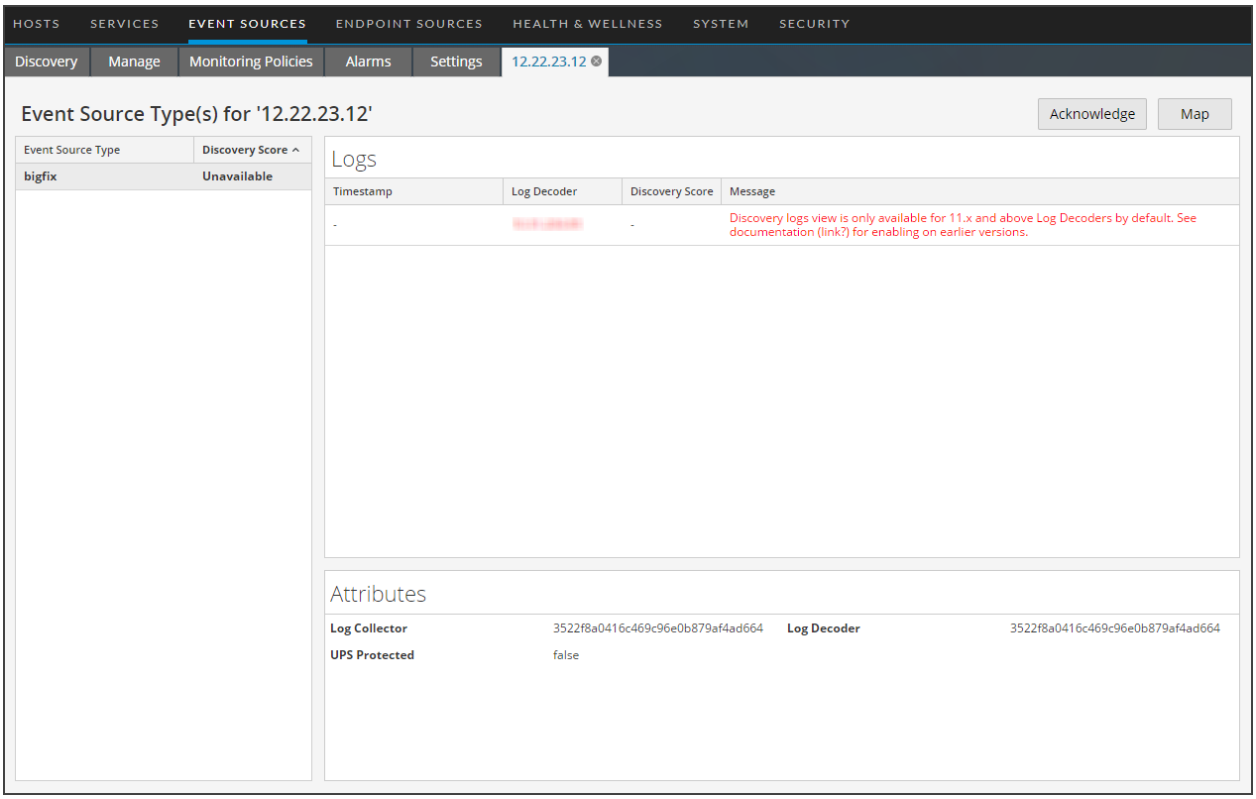
   Select Log Decoder service > **Explore** > **decoder** > **Properties** > **reset**. You select **reset** from a drop down menu.

Click **Send** after you select reset.

This is an example of the **index-logdecoder-custom.xml** file.

The following example shows the Discovery Score as **Unavailable** in the **Details** view for a Log Decoder.

The following example shows the message that is displayed in the Logs panel for a Log Decoder.


HOSTS
SERVICES
EVENT SOURCES
ENDPOINT SOURCES
HEALTH & WELLNESS
SYSTEM
SECURITY
Discovery
Manage
Monitoring Policies
Alarms
Settings
12.22.23.12
Event Source Type(s) for '12.22.23.12'
Acknowledge
Map
Event Source Type
Discovery Score
bigfix
Unavailable
Logs
Timestamp
Log Decoder
Discovery Score
Message
-
-
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
3522f8a0416c469c96e0b879af4ad664
Log Decoder
3522f8a0416c469c96e0b879af4ad664
UPS Protected
false

Upload a Log File to a Log Decoder

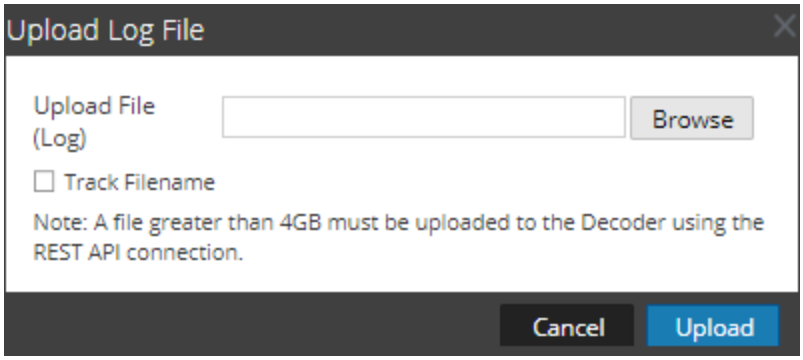
There are occasions when you want to analyze a log file that is not available on the service you are using. You can upload a log file captured on another service to NetWitness. Log filenames are of the type **.log**.

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 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. 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 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**.

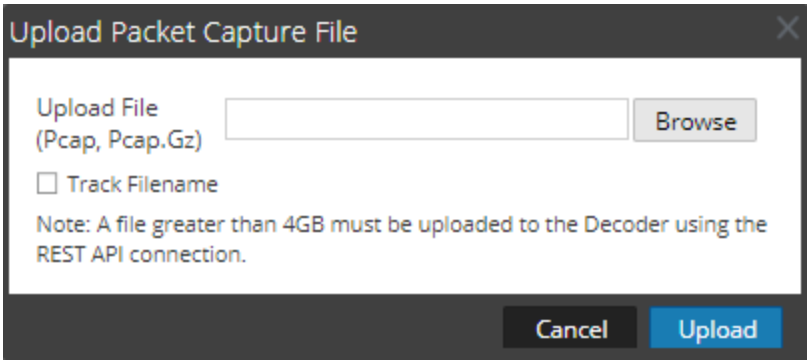
   The Administration Services view is displayed.

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.

F5 BIG IP - NetWitness Perfect Forward Secrecy Inspection Visibility

This topic discusses Perfect Forward Secrecy (PFS) Inspection Visibility with F5 BIG-IP Local Traffic Manager (LTM) and NetWitness Decoder integration.

NetWitness Decoder allows administrators to receive session keys from PFS sessions and decrypt each session traffic. It can receive session keys through its RESTful port, the PFS encrypted traffic through its capture interface, and supports the traffic decryption for inspection.

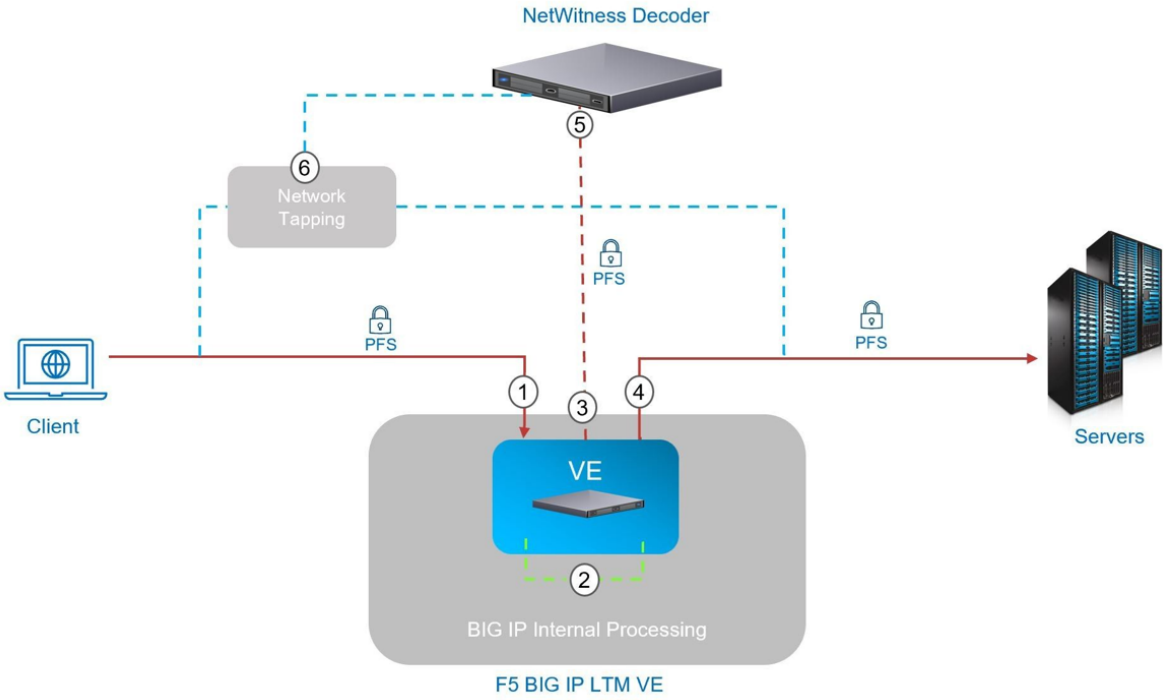
This solution utilizes functionality in iRule to extract session keys and deliver them to NetWitness Decoder via a sideband connection.

Prerequisites

A few prerequisites need to be considered in your environment for the integration of BIG-IP LTM and NetWitness Decoder.

- Deploy and Configure BIG IP LTM External Virtual Server (version 15.1.5 and above)
  - Configure Client SSL Profile with TLS 1.3 enabled.
  - Configure Server SSL Profile with TLS 1.3 enabled.
- Deploy and Configure NetWitness Decoder
  - Allow NetWitness Decoder RESTful port (default 50104) in your Network Firewall to receive keys from BIG-IP LTM.
    - NetWitness Decoder receives keys via HTTPS on RESTful port.
  - Create an Administrator account on the Decoder Device with role permissions `security.roles: decoder.manage` to forward TLS keys to Decoder.
    - Example name: bigip-service
  - Consider deploying NetWitness Decoder close to BIG-IP LTM so that it's feasible to mirror the encrypted traffic via TAP or Switch and avoid Network Latency in receiving TLS keys.

Deployment



1. A client connects to BIG-IP LTM using TLS with PFS cipher. The External Virtual Server has `clientssl` profile to decrypt traffic for inspection.
2. While inspecting the traffic an `iRule` on External Virtual Server collects the session keys and makes a `sideband` connection to Internal Virtual Server with the NetWitness Decoder as a `pool member`. The Internal Virtual Server has a `serverssl` profile to ensure traffic between BIG-IP LTM and NetWitness Decoder is protected since session keys are forwarded.
3. The Internal Virtual Server has a `HTTP profile` applied to manage the key transfer via HTTPS between BIG-IP LTM and NetWitness Decoder. `Oneconnect` profile is utilized and session keys will be sent on existing connections if available to reduce the connection setup overhead with NetWitness Decoder.
4. The External Virtual Server also has a `serverssl` profile to re-encrypt the traffic back to server pool.
5. The NetWitness Decoder receives the TLS keys from BIG-IP LTM on its RESTful port and decrypts the traffic ingested.
6. The Payload is not sent from BIG-IP LTM to NetWitness Decoder, this is received either from port mirroring configurations or a Network TAP infrastructure as shown above.

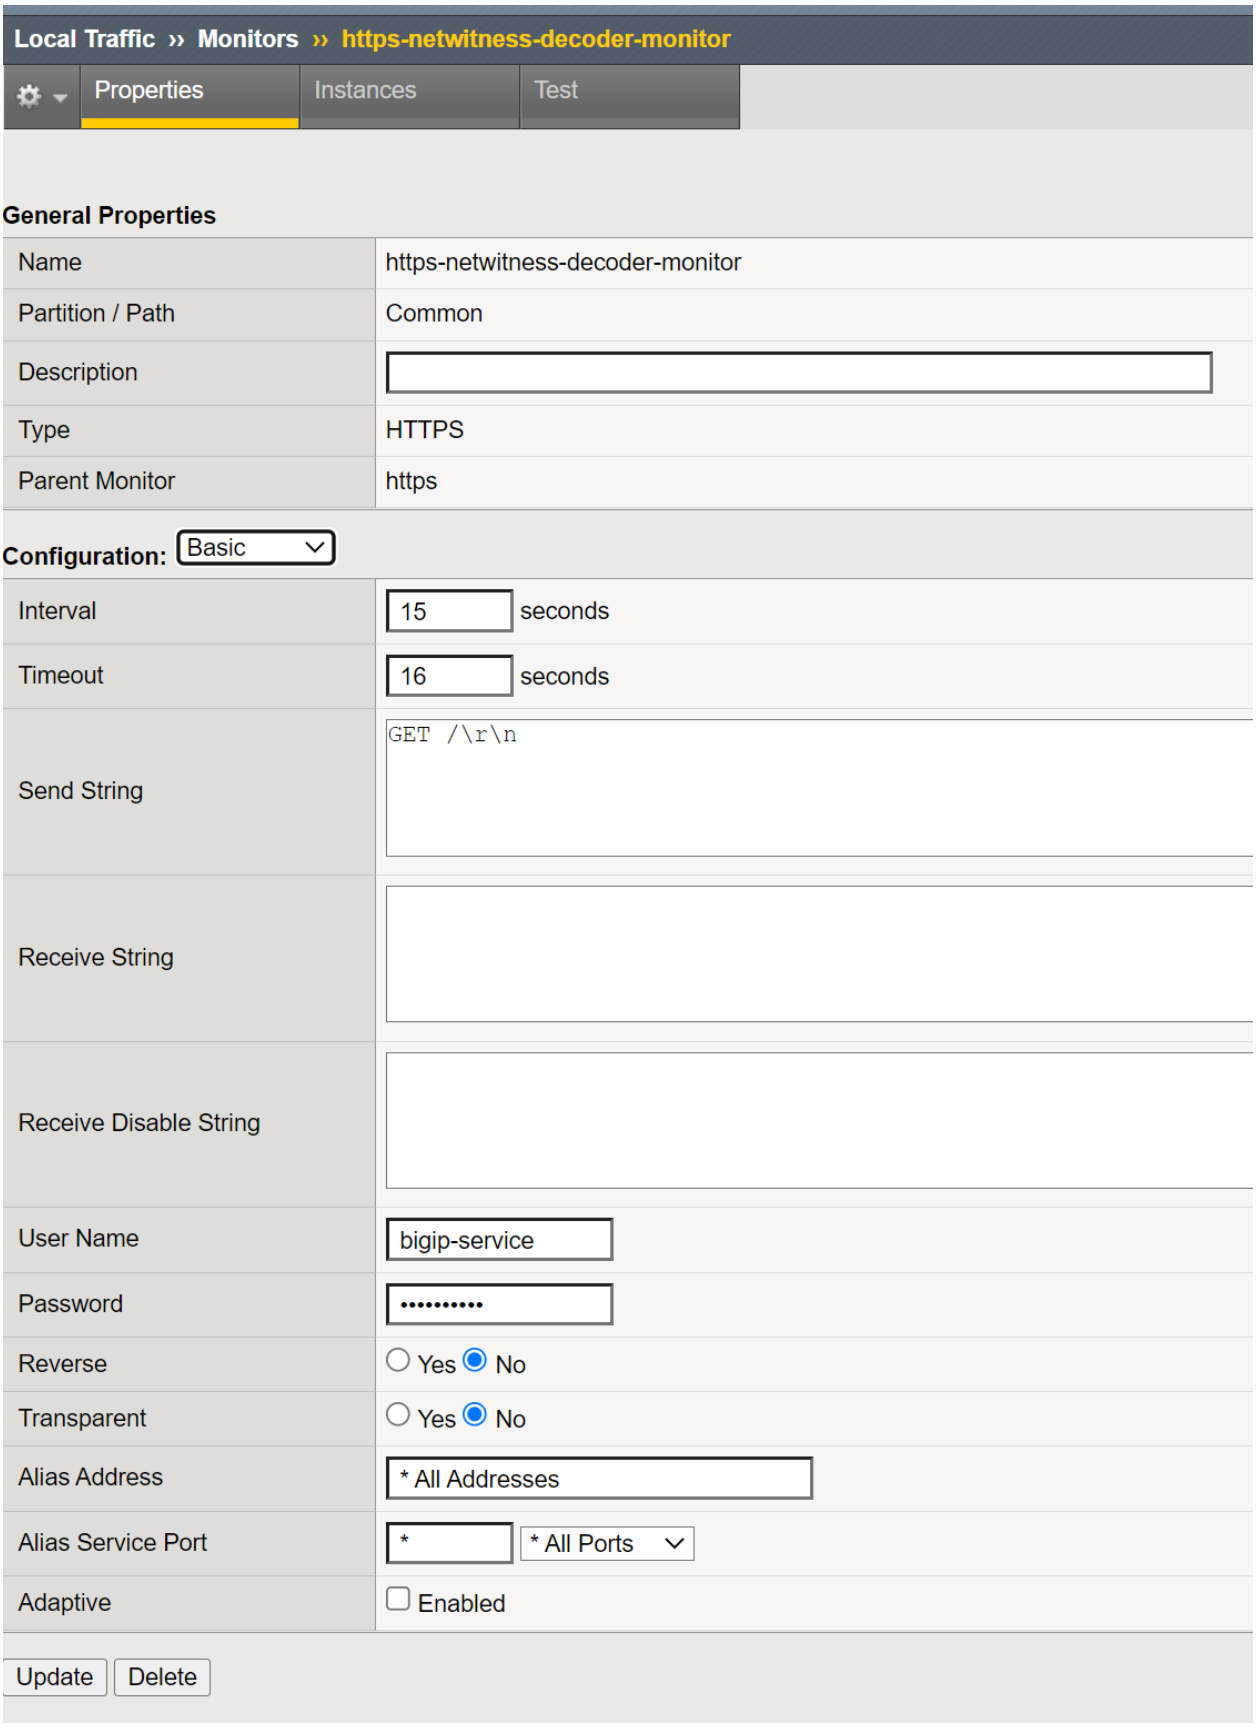
BIG-IP LTM Configurations

The BIG-IP configuration objects required include the following:

- HTTPS Monitor for monitoring NetWitness Decoder Node and Pool status (e.g online, offline etc…)
- Node and Pool objects for NetWitness Decoder.
- HTTP profile for managing key transfer to NetWitness Decoder via HTTPS.
- Server SSL profile for HTTPS sideband connection to NetWitness Decoder.
- Internal Virtual Server for receiving session keys and sideband communication with NetWitness Decoder.
- iRule: for session keys copying and sideband communication with Internal Virtual Server.

HTTPS Monitor

- Configure HTTPS monitor for NetWitness Decoder.
- Set the User Name and Password of the Administrator account created on Decoder device for BIG-IP transactions. Example Configuration using bigip-service account.

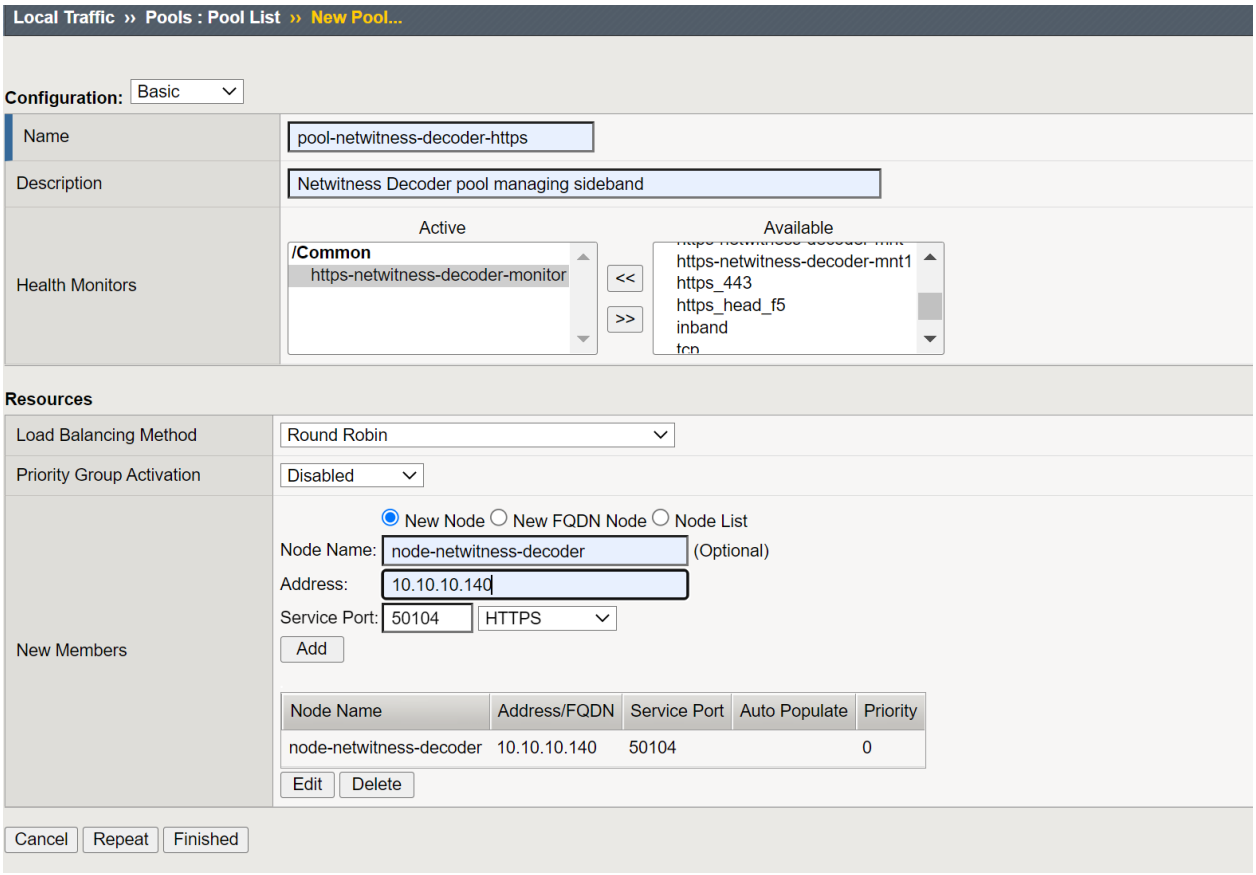

Local Traffic ›› Monitors ›› https-netwitness-decoder-monitor
Properties
Instances
Test
General Properties
Name
https-netwitness-decoder-monitor
Partition / Path
Common
Description
Type
HTTPS
Parent Monitor
https
Configuration: Basic
Interval
15
seconds
Timeout
16
seconds
Send String
GET /\r\n
Receive String
Receive Disable String
User Name
bigip-service
Password
Reverse
Yes
No
Transparent
Yes
No
Alias Address
* All Addresses
Alias Service Port
*
* All Ports
Adaptive
Enabled
Update
Delete

```
ltm monitor https /Common/https-netwitness-decoder-monitor
{
adaptive disabled
defaults-from /Common/https
destination *.50104
interval 15
ip-dscp 0
password <encrypted-bigip-service-account-password>
recv none
recv-disable none
send "GET /\r\n"
ssl-profile /Common/serverssl-secure
time-until-up 0
timeout 16
username bigip-service
}
```

Pool & Node

Configure Pool and Node objects.

- Select **Health Monitors** `https-NetWitness-decoder-monitor`
- Add **New Node** with the following:
  - Set **Node Name** example: `node-NetWitness-decoder`
  - **Address**: <Decoder IP Address >, **Service Port**: 50104 HTTPS and then click **Add**.
- Click **Finished** to create Pool with a Node.




```
ltm pool /Common/pool-netwitness-decoder-https {
members {
/Common/node-netwitness-decoder:50104 {
address 10.10.10.140
}
}
monitor /Common/https-netwitness-decoder-monitor
}
```

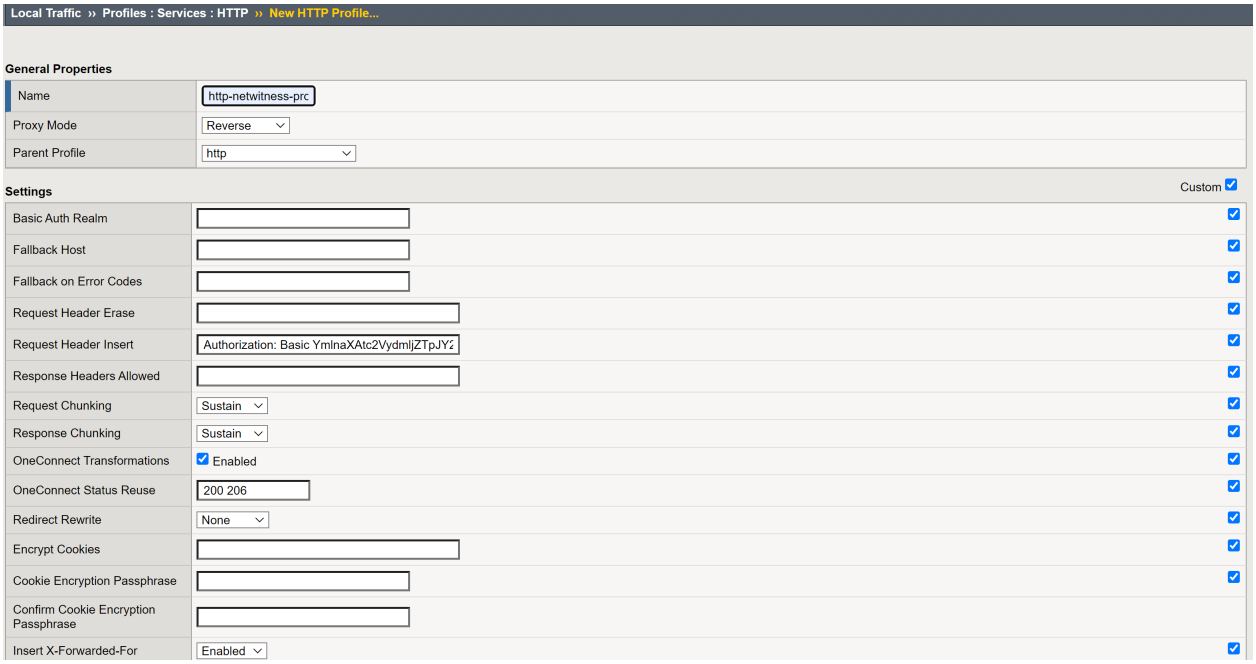
HTTP Profile

- Configure HTTP Profile **Name** `http-NetWitness-profile` to authenticate NetWitness Decoder device and send TLS keys.
- Set **Parent Profile** as **http**
- Check the **Custom** check box to edit the Settings.
- Set **Request Header Insert** field with basic Authorization header string
  - `"Authorization: Basic <Base64 encoded stirng (username:password)>"`
  - Example: `"Authorization: Basic YmlnaXAtc2VydmljZTpJY2FuU2VleW91QDEyMyE="`

  - Where
    - username - Administrator account on Decoder device (example: bigip-service)
    - password - The account password

- Set **Insert X-Forwarded-For** as **Enabled** , this would add X-Forwarded-For header with the original Client IP of the session.

Rest of the other configurations can remain as default values.

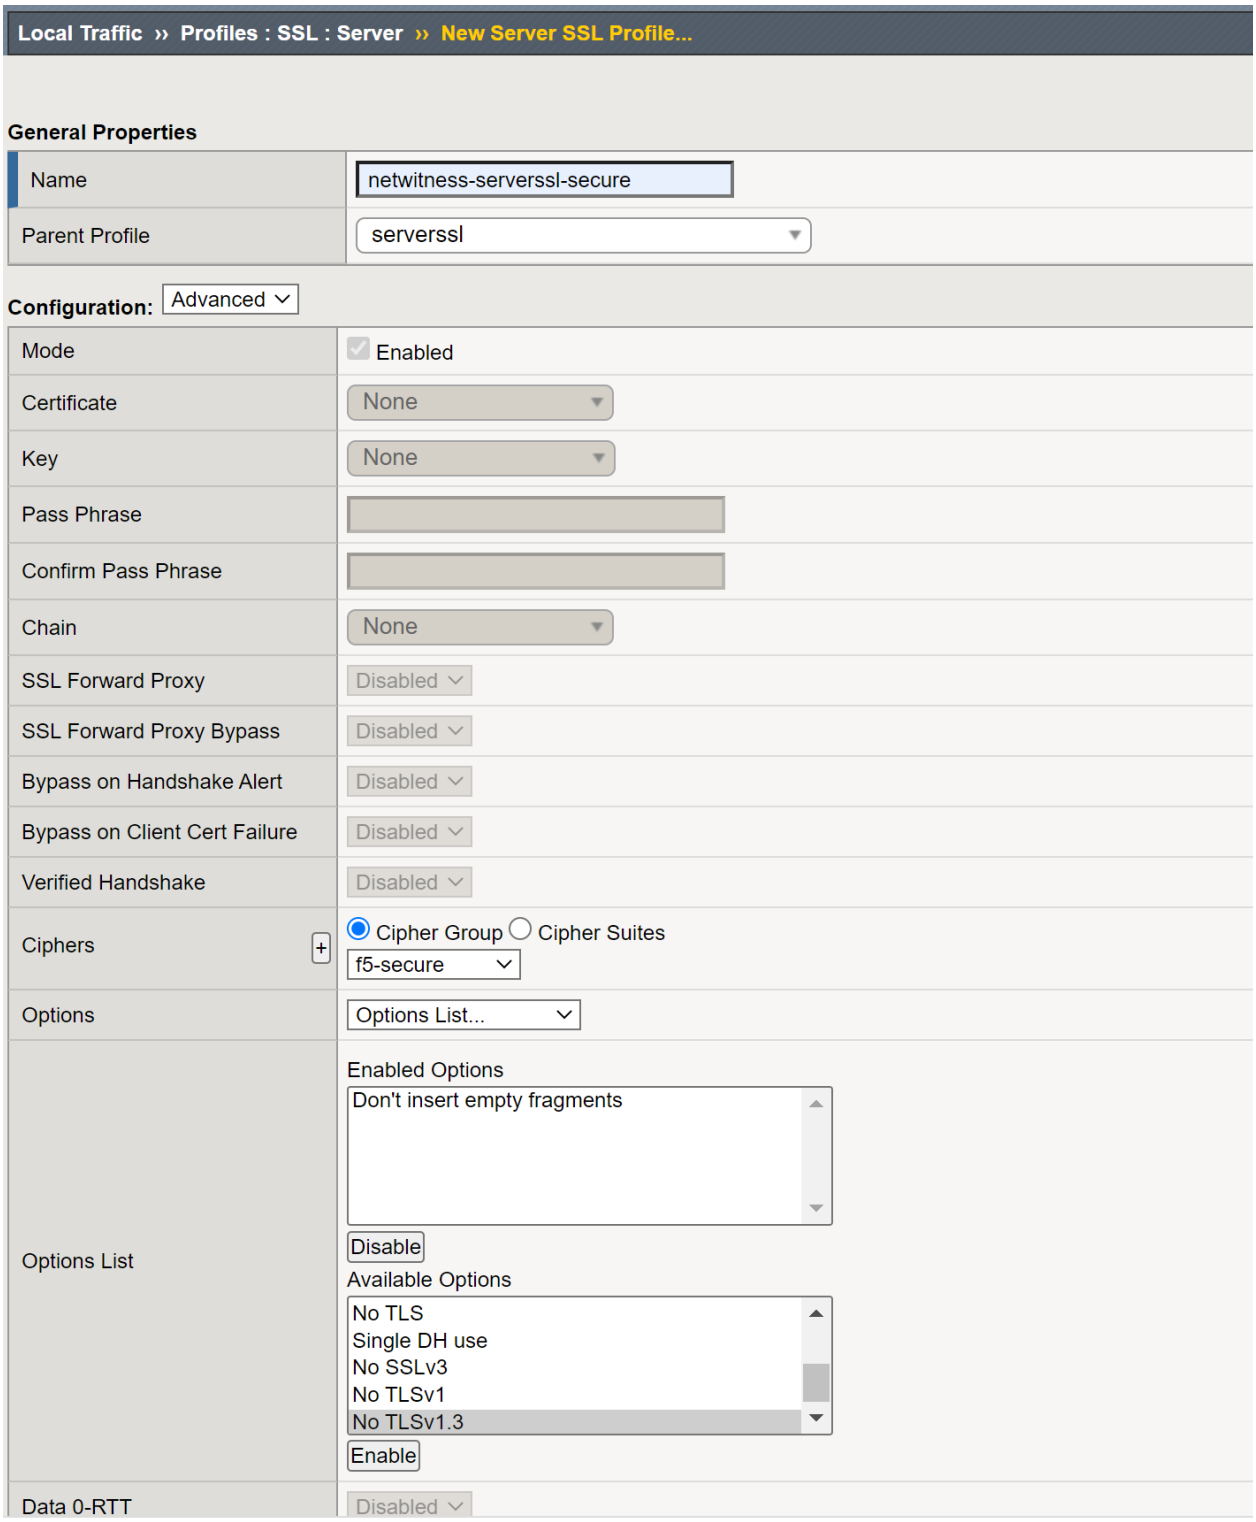



```
ltm profile http /Common/http-NetWitness-profile {
accept-xff disabled
app-service none
basic-auth-realm none
defaults-from /Common/http
encrypt-cookies none
enforcement {
known-methods { CONNECT DELETE GET HEAD LOCK OPTIONS POST PROPFIND PUT
TRACE UNLOCK }
max-header-count 64
max-header-size 32768
max-requests 0
pipeline allow
rfc-compliance disabled
truncated-redirects disabled
unknown-method allow
}
fallback-host none
```

```
fallback-status-codes none
header-erase none
header-insert "Authorization: Basic
YmlnaXAtc2VydmljZTpJY2FuU2VleW91QDEyMyE="
hsts {
include-subdomains enabled
maximum-age 16070400
mode disabled
preload disabled
}
insert-xforwarded-for enabled
lws-separator none
lws-width 80
oneconnect-status-reuse "200 206"
oneconnect-transformations enabled
proxy-type reverse
redirect-rewrite none
request-chunking sustain
response-chunking sustain
response-headers-permitted none
server-agent-name BigIP
sflow {
poll-interval 0
sampling-rate 0
}
via-request preserve
via-response preserve
xff-alternative-names none
}
```

Server SSL Profile

1. Set server-ssl profile **Name** `NetWitness-serverssl-secure`
2. Select **Advanced** from **Configuration**.
3. Set **f5-secure** from **Ciphers > Cipher Group**.
4. (Optional) **Disable** `no-TLS1.3` in options which enables TLS 1.3 between BIG-IP and NetWitness Decoder.
5. Click **Update** to create `server-ssl` profile.

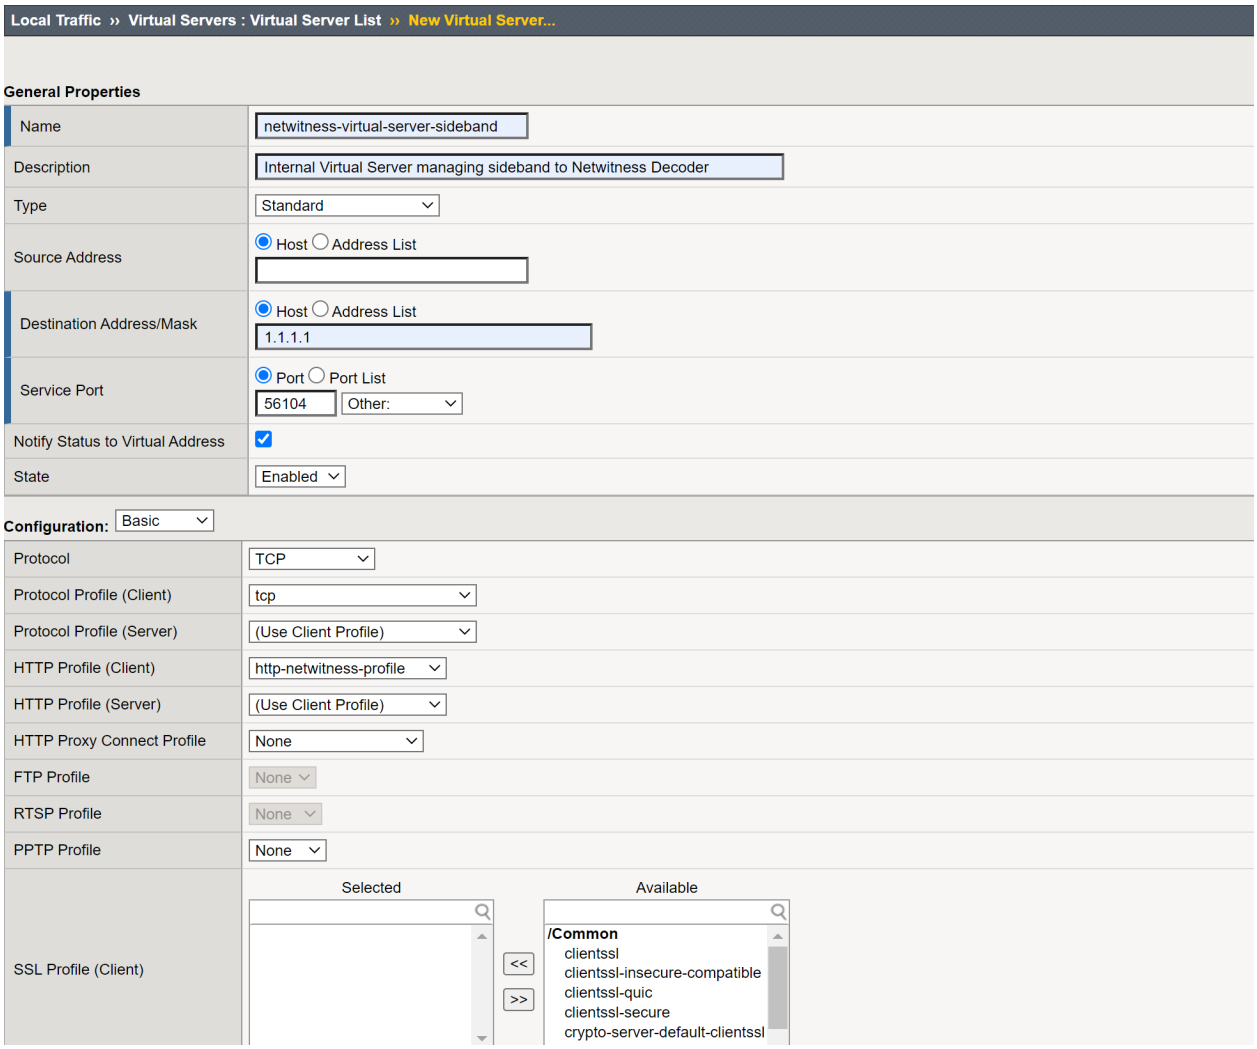



```
ltm profile server-ssl /Common/NetWitness-serverssl-secure {
app-service none
cipher-group /Common/f5-secure
ciphers none
```

```
defaults-from /Common/serverssl
options { no-ssl }
renegotiation disabled
}
```

Internal Virtual Server

1. Create an Internal Virtual server to manage sideband connection to NetWitness Decoder and forward keys.
2. Set **Name** as `NetWitness-virtual-server-sideband` and add **Description** for virtual server.
3. Let the Source Address be blank, and it would default to 0.0.0.0.
4. Set **Destination Address/Mask** to 1.1.1.1; this is a non-routable IP address and is used internally.
5. Set **Service Port** to 56104 , a virtual non-routable port used internally.
6. Select **HTTP Profile (Client)** `http-NetWitness-profile` that was created earlier. By default, Server inherit Client profile.
7. Select **SSL Profile (Server)** `NetWitness-serverssl-secure` which was created earlier. This would be used by the Pool `pool-NetWitness-decoder-https` when connecting to NetWitness Decoder.
8. Select **Source Address Translation** Auto-Map.
9. Select **OneConnect Profile** oneconnect.
10. Select **Default Pool** `pool-NetWitness-decoder-https` which was created earlier.
11. Click **Finished** to create Internal Virtual Server.




```
ltm virtual /Common/NetWitness-virtual-server-sideband {
description "Internal Virtual Server managing sideband to NetWitness
Decoder"
destination /Common/1.1.1.1:56104
ip-protocol tcp
mask 255.255.255.255
pool /Common/pool-NetWitness-decoder-https
profiles {
/Common/http-NetWitness-profile { }
/Common/NetWitness-serverssl-secure {
context serverside
}
/Common/oneconnect { }
/Common/tcp { }
}
```

```
serverssl-use-sni disabled
source 0.0.0.0/0
source-address-translation {
type automap
}
translate-address enabled
translate-port enabled
}
```

iRule Session keys copying and Sideband communication

Configuration

1. Configure and add the following `netwitness-sideband-irule` **iRule** in your **External Virtual Server**, as an Administrator
2. Set **variable** `static::virtual_server` in `RULE_INIT` with name of **Internal Virtual Server** of your environment.
   - Example `set static::virtual_server "NetWitness-virtual-server-sideband"`
3. (Optional) If you would like to disable extracting keys for Client side communication, then comment the routine CLIENTSSL_HANDSHAKE.
4. (Optional) If you would like to disable extracting keys for Server side communication, then comment the routine SERVERSSL_HANDSHAKE.
5. (Optional) For any additional troubleshooting uncomment the log messages.

iRule `NetWitness-sideband-irule`

```
# NetWitness sideband iRule.


# The script extracts TLS keys from PFS (Perfect Forward Secrecy) communication between client and server, then it forwards
# the keys to internal virtual server via sideband connection. The internal virtual server then send the keys to
# NetWitness Decoder for TLS Decryption.
#
# Configuration Steps:
#
#   1. (Required) As an administrator you would need to set variable static::virtual_server with name of internal virtual server of your environment.
#   2. (Optional) If you would like to disable extracting keys for Client side communication, then comment the routine CLIENTSSL_HANDSHAKE.
```

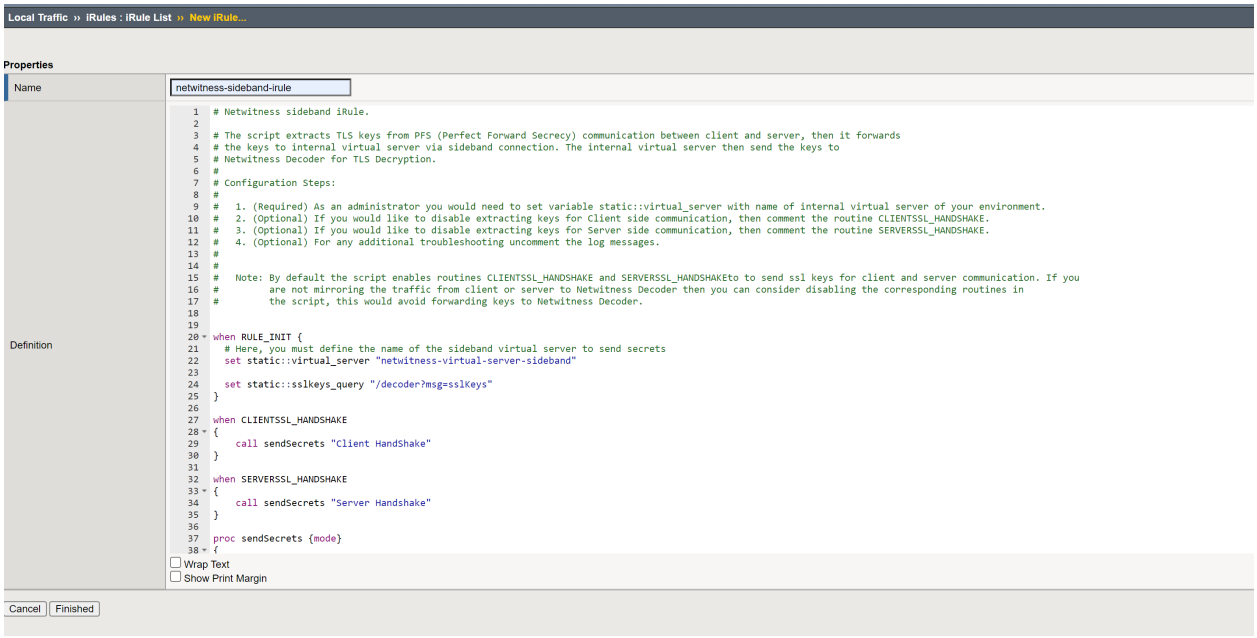
```
# 3. (Optional) If you would like to disable extracting keys for Server side communication, then comment the routine SERVERSSL_HANDSHAKE.
# 4. (Optional) For any additional troubleshooting uncomment the log messages.
#
#
# Note: By default the script enables routines CLIENTSSL_HANDSHAKE and SERVERSSL_ HANDSHAKEto to send ssl keys for client and server communication. If you
#       are not mirroring the traffic from client or server to NetWitness Decoder then you can consider disabling the corresponding routines in
#       the script, this would avoid forwarding keys to NetWitness Decoder.
when RULE_INIT {
# Here, you must define the name of the sideband virtual server to send secrets
set static::virtual_server "NetWitness-virtual-server-sideband"
set static::sslkeys_query "/decoder?msg=sslKeys"
}
when CLIENTSSL_HANDSHAKE
{
call sendSecrets "Client HandShake"
}
when SERVERSSL_HANDSHAKE
{
call sendSecrets "Server Handshake"
}
proc sendSecrets {mode}
{
if {[SSL::cipher version] == "TLSv1.3" } {
#log local0. "$mode: CLIENT_HANDSHAKE_TRAFFIC_SECRET [SSL::clientrandom] [SSL::tls13_ secret client hs]"
set secrets "random=[SSL::clientrandom]"
if { [SSL::tls13_secret client early] ne "" } {
set secrets "$secrets&CETS=[SSL::tls13_secret client early]"
}
set secrets "$secrets&CHTS=[SSL::tls13_secret client hs]&SHTS=[SSL::tls13_secret server hs]&CTS0=[SSL::tls13_secret client app]&STS0=[SSL::tls13_secret server app]"
if { [catch {call sidebandSend $secrets} ] } {
#log local0. "$mode $static::virtual_server sideband is not reachable, ssl version [SSL::cipher version]"
return
}
```

```
}
else
{
set secrets "random=[SSL::clientrandom]&premaster=[SSL::sessionsecret]"
if { [catch {call sidebandSend $secrets} ] } {
#log local0. "$mode $static::virtual_server sideband is not reachable, ssl version [SSL::cipher version]"
return
}
}
}
proc sidebandSend {secrets} {
#log local0. "$static::virtual_server launching sidebandSend ..."
set cmp_unit [TMM::cmp_unit]
# Key for session table
set key "${cmp_unit}_conn_${static::virtual_server}"
# Session table data for $key for this dest_addr
set conn [session lookup dest_addr $key]
if { $conn eq "" } {
set conn [connect -timeout 1000 -idle 300 -status conn_status $static::virtual_server]
if { $conn_status ne "connected" }{
#log local0. "Failed to connect to virtual_server $static::virtual_server"
return
}
session add dest_addr $key "$conn" 300
}
else {
# Attempt sideband connection re-use
set conn_info [connect info -status $conn]
set conn_state [lindex [lindex $conn_info 0] 0]
if { $conn_state ne "connected" }{
set conn [connect -timeout 1000 -idle 300 -status conn_status $static::virtual_server]
if { $conn_status ne "connected" }{
#log local0. "Failed to connect on connection re-use to virtual_server $static::virtual_server"
return
}
session add dest_addr $key "$conn" 300
```

```
}
}
#log local0. "Sending keys to virtual_server $static::virtual_server ..."
set send_bytes [send -timeout 1000 -status send_status $conn "GET $static::sslkeys_query&$secrets HTTP/1.0\r\n\r\n"]
#log local0. "Sent keys to virtual_server $static::virtual_server ..."
recv -timeout 1 $conn
}
```

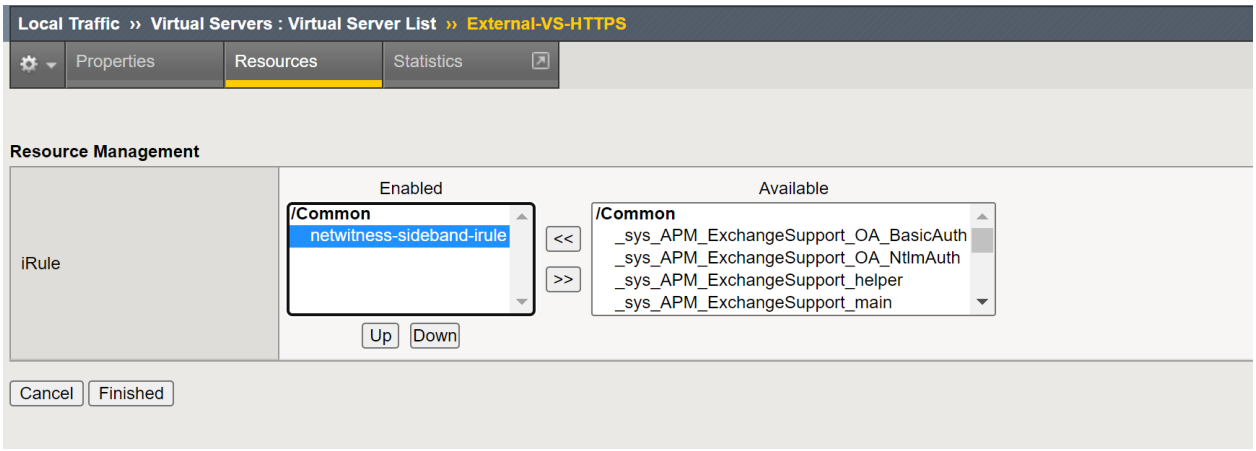
Copy

1. **Create iRule** in iRule List create wizard
2. Copy the above iRule content to **Definition**
3. Click **Finished** to save iRule



Example: Add iRule to External Virtual Server

Select **iRule** `netwitness-sideband-irule` and click **Finished** to add iRule.


Local Traffic ›› Virtual Servers : Virtual Server List ›› External-VS-HTTPS
Properties
Resources
Statistics
Resource Management
iRule
Enabled
/Common
netwitness-sideband-irule
Up
Down
<<
>>
Available
/Common
_sys_APM_ExchangeSupport_OA_BasicAuth
_sys_APM_ExchangeSupport_OA_NtlmAuth
_sys_APM_ExchangeSupport_helper
_sys_APM_ExchangeSupport_main
Cancel
Finished

Troubleshooting Packet Drops

Packet drops can occur if there is backup happening at Packet Pool or Session Pool on the Decoder, and eventually, Decoder runs out of free capture pages resulting in drops.

Quick configuration Checks to avoid packet drops

The Decoder would log warning messages when it encounters packet drops. These logs contain possible reasons for drops, and you can solve some of the drop symptoms through simple configuration checks.

To check and tune the configuration:

In most cases, the decoder configuration parameters would be different from the default configuration or with the hardware deployed. So, make sure to check the following and fix configuration issues.

1. When `/var/log/messages` logs show drops with reason "check capture configuration, packet sizes".
   - NwDecoder[15913]: [Packet] [warning] **Packet drops encountered, packet capture** (737626/737628): **check capture configuration, packet sizes and rates**
   - For 10G decoder run

     `/decoder reconfig op=10g update=1`

     Or

     Explore view | /decoder | properties | select reconfig | op=10g update=1 | send
   - For normal decoder run

     `/decoder reconfig update=1`

     Or

     Explore view | /decoder | properties | select reconfig | update=1 | send
   - The Decoder service needs to be restarted for changes to be effective.
   - Monitor decoder for drops
2. When `/var/log/messages` logs show drops with reason "check packet database configuration, iostats, packet and content calls"
   - NwDecoder[74030]: [Packet] [warning] **Packet drops encountered, packet write** (717957/723314): **check packet database configuration, iostats, packet and content calls**
   - Packet drops encountered (884632/884642): **check session & meta database configuration, iostats and sdk activity**.
   - For 10G decoder run
     - `/database reconfig op=10g update=1`

       Or

- Explore view | /database | properties | select reconfig | op=10g update=1 | send

- For normal decoder run
  - `/database reconfig update=1`

    Or
  - Explore view | /database | properties | select reconfig | update=1 | send
- The Decoder service needs to be restarted for changes to be effective.
- If we are still seeing drops, then a few more changes are required. There could be i/o bound waits for database writes on the Decoder.
  - Set `/database/config/packet.integrity.flush=normal`
  - Set `/database/config/session.integrity.flush=normal`
  - Set `/database/config/meta.integrity.flush=normal`
- Monitor decoder for drops.

3. As of NetWitness Platform latest version, `sosreport` retrieves service and database reconfig information as well as what settings were active when the `sosreport` was retrieved. This information can be used for cross-checking purposes and can be found in the `..../sos_commands/rsa_nw_rest` directory (service-reconfig, database-reconfig, ls<service>).

Information required to troubleshoot packet drops

- Monitor Packet Drops Tool Output (Highly recommended)
- This can be accessed through REST port `http://<decoder>:50104/sdk/app/packetdrops`
- By default, the tool looks for drops in last 24 hrs and also provides options to search drops based on time ranges.
- Enable detailed stats on decoder `REST /decoder/parsers/config/detailed.stats=yes`

What if the REST port is inaccessible?

- Enable detailed stats on decoder

  `/decoder/parsers/config/detailed.stats=yes`
- Wait for new drops and then collect `sosreport`. The `sosreport` on Decoder would collect few stats db files.
- These files can be copied to your local decoder `/var/netwitness/decoder/statdb` and restart Decoder.
- Access packet drops tool on your local Decoder.

How do I troubleshoot packet drops?

The packet drops tool would help narrow down the possible cause for drops. There could be various cases involved that need verification.

The following sections describe how to use the Packet drops tool and analyze its results.

Navigate to the Packet drops tool using the REST interface

- The tool would search drop instances and list them with links.
- Start investigating the latest drop instances where the drop count is high.
- Look at other drop instances and find the pattern for drops.

Introduction to Drops tool Charts

- Incremental Packet Drops - Displays packet drops count at that instance of time
- Capture - Displays traffic ingestion rates Capture.rate in Mbps and Calculated capture rate (instantaneous rate) in Mbps.
- Packet Pool - Displays graphs for the following packet pool stats:
  - capture - /decoder/stats/pool.packet.capture → Number of free packet pages available for capture
  - assembler - /decoder/stats/pool.packet.assembler → Number of packet pages waiting to be assembled.
  - write - /decoder/stats/pool.packet.write → Number of packet pages waiting to be updated to Packet Database.
  - pool - /decoder/stats/assembler.packet.pages → Number of packet pages held in the assembler
  - export - /decoder/stats/pool.packet.export → Number of packet pages waiting to be exported
- Session Pool
  - This displays graphs for following session pool stats
  - parse - /decoder/parsers/stats/pool.session.parse → Number sessions waiting to be parsed.
  - write - /decoder/stats/pool.session.write → Number of session pages waiting to be written.
  - correlate - /decoder/stats/pool.session.correlate → Number of session pages waiting to be correlated.
  - queue.sessions.total - /decoder/parsers/stats/queue.sessions.total → The total number sessions in parse threads and queues.
  - export - /decoder/stats/pool.session.export Number of session pages waiting to be exported.

When packet drops occur, the chart shows that instance of time in yellow.

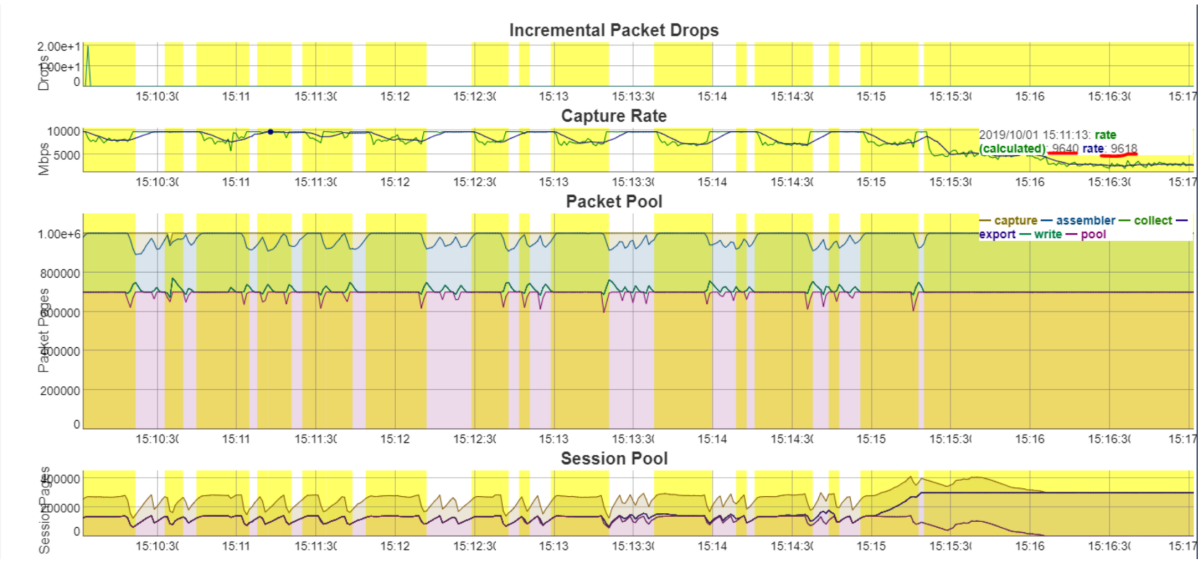
Symptoms Checklist

We recommend following the symptom checks in the order listed below, and this helps to troubleshoot all the symptoms associated with current packet drops.

Symptom 1: Higher traffic ingestion rates for the content deployed would cause packet drops

If the traffic ingestion rate -(capture) > 8 Gbps for the content (higher number of parsers) deployed on Decoder, this would cause packet drops.

ex: Below screenshot shows capture rate > 8 Gbps (9.6 Gbps)



Resolution:

Consider the following options:

- Split the traffic ingestion into multiple decoders so that the ingestion rate would be ~4-5 Gbps

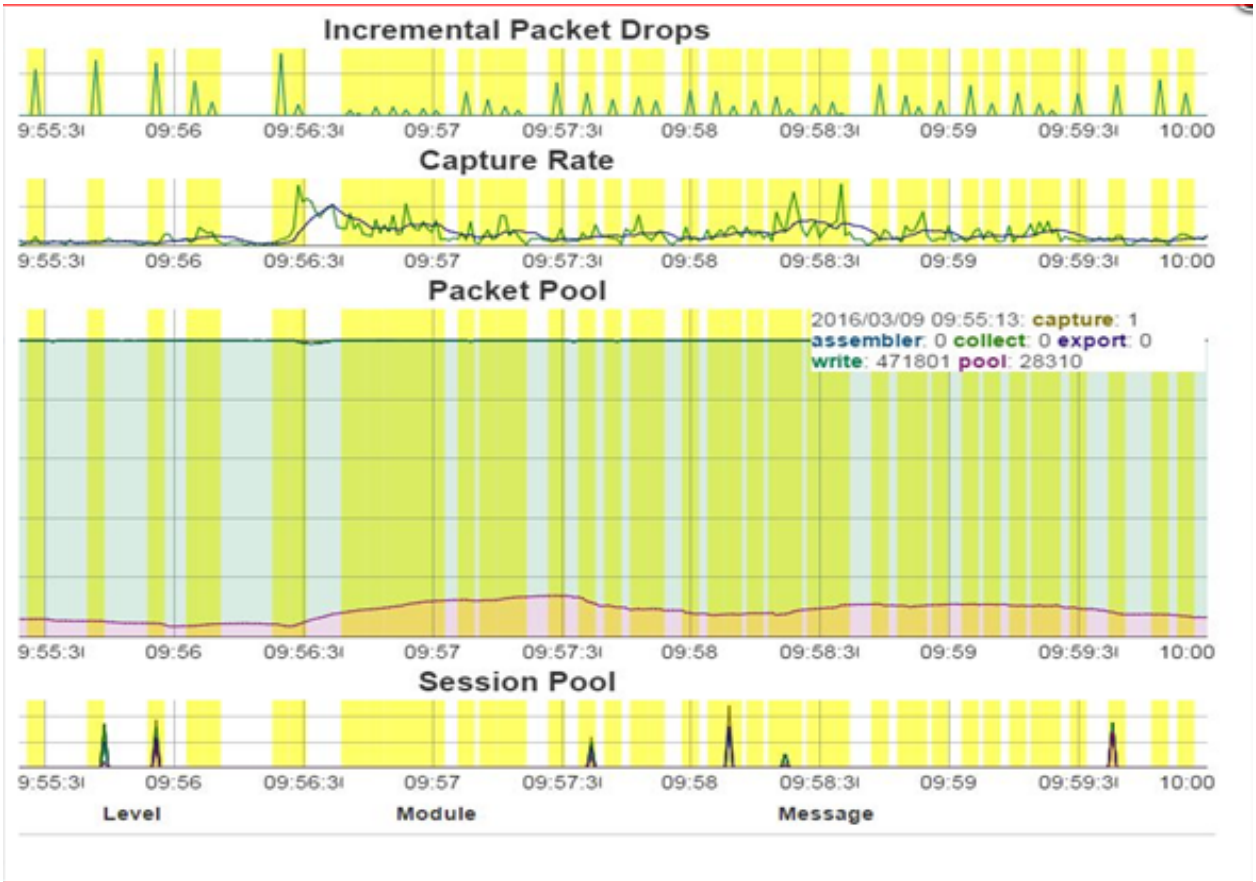
  Or

- Start with baseline supported 10G content and then add parsers one by one until no drops are observed.
  - Refer Parsing and Content Considerations section on Configure High Speed Packet Capture Capability

Symptom 2: Packet Database Write backup would cause packet drops

Packet Pool Write - write (pool.packet.write), if write backup increases, then packet write delays would be causing the drops

ex: Below screenshot displays, there are many packet pages (471K) waiting on the write queue.



Decoder logs would throw warnings like the below:

- NwDecoder[74030]: [Packet] [warning] **Packet drops encountered, packet write** (717957/723314): **check packet database configuration, iostats, packet and content calls**

Top -Hp <Decoder PID> would display decoder packet write thread waiting on Disk.

```
%Cpu(s):  0.9 us,  0.7 sy,  0.0 ni, 90.3 id,  8.1 wa,  0.0 hi,  0.0 si,  0.0 st
KiB Mem : 13148560+total, 27355252 free, 72468064 used, 31662292 buff/cache
KiB Swap:  4194300 total,  4190452 free,     3848 used. 58113828 avail Mem

  PID USER      PR  NI    VIRT    RES    SHR S  %CPU %MEM     TIME+ COMMAND
131928 root      20   0   45.8g   4.6g 545220 S   9.0  3.6 452:10.36 PFRING Balancer
131851 root      20   0   45.8g   4.6g 545220 S   1.7  3.6 723:05.92 Parse Work 2
131859 root      20   0   45.8g   4.6g 545220 S   1.7  3.6 723:46.83 Parse Work 10
131850 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 722:14.60 Parse Work 1
131853 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 724:39.44 Parse Work 4
131855 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 722:58.13 Parse Work 6
131857 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 724:08.50 Parse Work 8
131858 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 724:45.26 Parse Work 9
131860 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 723:19.92 Parse Work 11
131861 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 722:48.85 Parse Work 12
131862 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 723:51.35 Parse Work 13
131863 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 723:17.89 Parse Work 14
131852 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 723:02.83 Parse Work 3
131854 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 722:38.98 Parse Work 5
131856 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 722:40.62 Parse Work 7
131864 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 721:45.72 Parse Work 15
131865 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 725:29.23 Parse Work 16
131838 root      20   0   45.8g   4.6g 545220 S   0.7  3.6 917:20.65 CaptureWork
131839 root      20   0   45.8g   4.6g 545220 S   0.7  3.6 526:43.29 Assem Work
131840 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  28:04.24 Packet Export W
131841 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  42:24.57 Collection ID W
131843 root      20   0   45.8g   4.6g 545220 D   0.7  3.6 303:38.71 Session Write W
131844 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  41:22.46 Session Id Work
131845 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  41:23.96 Session Correla
131849 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  41:19.76 Session Export
132932 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  49:56.33 Request Handler
131771 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  31:57.08 stat updates
131842 root      20   0   45.8g   4.6g 545220 D   0.3  3.6 846:47.73 Packet Write Wo
131847 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  61:44.11 Session Index W
131848 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  38:53.76 Session Stream
131925 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  23:09.92 Timestamper
138159 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  48:02.02 Request Handler
131770 root      20   0   45.8g   4.6g 545220 S   0.0  3.6  20:36.10 NwDecoder
131773 root      20   0   45.8g   4.6g 545220 S   0.0  3.6  20:24.78 service listen
131790 root      20   0   45.8g   4.6g 545220 S   0.0  3.6  51:13.34 Request Handler
131798 root      20   0   45.8g   4.6g 545220 S   0.0  3.6   0:01.70 REST executor
131837 root      20   0   45.8g   4.6g 545220 S   0.0  3.6   0:03.07 Request Handler
131846 root      20   0   45.8g   4.6g 545220 S   0.0  3.6   0:00.10 FeedLoader
131868 root      20   0   45.8g   4.6g 545220 S   0.0  3.6   2:18.88 amqp thread
```

iotop tool can show Disk i/o activity: rpm is available here and can be installed: `http://mirror.centos.org/centos/7/os/x86_64/Packages/iotop-0.6-4.el7.noarch.rpm`

Example iotop results where decoder packet write thread is blocked on IO 99% of decoder IO time and its throughput is just ~400KB/s

```
iotop -o -d 2
```

```
Total DISK READ :       0.00 B/s | Total DISK WRITE :     463.05 K/s
Actual DISK READ:       0.00 B/s | Actual DISK WRITE:       0.00 B/s
  TID  PRIO  USER     DISK READ  DISK WRITE  SWAPIN     IO>    COMMAND
 3373 be/4 root        0.00 B/s  431.79 K/s  0.00 % 99.99 % NwDecoder [Packet Write Wo]
 3374 be/4 root        0.00 B/s   31.26 K/s  0.00 % 99.99 % NwDecoder [Session Write W]
```

Resolution:

- Make sure database configuration tuning is applied as suggested in the section To check and tune the configuration:

- Check I/O stats on the Decoder using the command "`iostat -mNx 1`". Refer to How do you get statistics on I/O performance?
  - If % `iowait` is > 10% then decoder packet db writes have higher i/o waits
  - If % util for packetdb goes greater than 95 and wMB/s < 1000 , then Disk write throughput is low and the Disk where packetdb exists needs to be replaced.
  - If `iotop` is installed the disk io activity can be monitored through 'iotop –o –d 2'
  - For a 10G decoder we recommend packet db disk write throughput to be 1300 MB/s (~10Gbps) for better write performance.
- Lot of Content calls to extract packets or content can cause packet write issue.
  - Check drops tool logs or `/var/log/messages` or sosreport logs for SDK-Content Calls.
  - You can also use NwConsole **topQuery** command on messages logs to identify Content calls.
  - Set `/decoder/sdk/config/packet.read.throttle=100` (a higher value) so that packet write would get preference.
  - Check service invoking SDK-Content calls and reduce the content calls.
- Kernel and Driver compatibility issues
  - Check if the firmware is updated according to Kernel version. If not update firmware.

How do you get statistics on I/O performance?

1. The command you want to run for near real time statistics on I/O usage is "`iostat -N -x -m 1`". For detailed information on the output of iostat, type "`man iostat`". If the columns do not line up, you can leave off the `-N` option, but you should probably run it once so you can see what disk groups correspond to which databases.

	%user	%nice	%system	%iowait	%steal	%idle
avg-cpu:	2.41	0.00	0.22	5.98	0.00	91.39

Device	rrqm/s	wrqm/s	r/s	w/s	rMB/s	wMB/s	avgrq-sz	avgqu-sz	await	svctm	%util
sda	0.00	0.00	0.00	1.00	0.00	0.00	8.00	0.00	3.00	3.00	0.30
sdb	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_root	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

VolGroup-lv_home	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_swap	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_nwhome	0.00	0.00	0.00	1.00	0.00	0.00	8.00	0.00	3.00	3.00	0.30
VolGroup-lv_tmp	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_vartmp	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_varlog	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
sdc	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
sdd	0.00	0.00	242.00	0.00	30.25	0.00	255.97	1.98	8.14	4.13	100.00
sde	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
index-index	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
concentrator-root	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
concentrator-sessiondb	0.00	0.00	141.00	0.00	17.62	0.00	255.94	0.81	5.84	5.77	81.30
concentrator-metadb	0.00	0.00	102.00	0.00	12.75	0.00	256.00	1.17	11.39	9.75	99.40

2. In the above sample output, you can see the session and meta DBs (last two lines) are reading in about 30 MB/s combined. You can also see that `sdd` corresponds to those two databases, since the numbers combined roughly add up. The important takeaway is the last column, %util. This means: Percentage of CPU time during which I/O requests were issued to the device (bandwidth utilization for the device). Device saturation occurs when this value is close to 100%. Therefore, the disks are currently saturated with read requests and the databases are currently I/O bound. When measuring performance of the SA Core services, it's important to know if the software is I/O bound. This would be the limiting factor for increasing performance as the current hardware is running at capacity.
3. To determine which physical volume the index is on, run "`vgdisplay -v index`". In the above example, it's /dev/sde. To see a list of volumes, run "`vgdisplay -s`".
4. Run "`ls -l /dev/mapper`" to see the drive mappings.

Symptom: Session Database Write backup would cause packet drops

Session Pool Write - write (pool.session.write), if session write backup increases, then session write delays would be causing the drops

ex: Below screenshot displays there are many session pages (198K) waiting on session write queue



Decoder logs would also throw warnings like below:

- Packet drops encountered (884632/884642): **check session & meta database configuration, iostats and sdk activity**.

Top -Hp <Decoder PID> would display decoder packet write thread waiting on Disk.

```
Cpu(s):  0.9 us,  0.7 sy,  0.0 ni, 90.3 id,  8.1 wa,  0.0 hi,  0.0 si,  0.0 st
iB Mem : 13148560+total, 27355252 free, 72468064 used, 31662292 buff/cache
iB Swap:  4194300 total,  4190452 free,     3848 used. 58113828 avail Mem

  PID USER      PR  NI    VIRT    RES    SHR S %CPU %MEM     TIME+ COMMAND
31928 root      20   0   45.8g   4.6g 545220 S  9.0  3.6 452:10.36 PFRING Balancer
31851 root      20   0   45.8g   4.6g 545220 S  1.7  3.6 723:05.92 Parse Work 2
31859 root      20   0   45.8g   4.6g 545220 S  1.7  3.6 723:46.83 Parse Work 10
31850 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 722:14.60 Parse Work 1
31853 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 724:39.44 Parse Work 4
31855 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 722:58.13 Parse Work 6
31857 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 724:08.50 Parse Work 8
31858 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 724:45.26 Parse Work 9
31860 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 723:19.92 Parse Work 11
31861 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 722:48.85 Parse Work 12
31862 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 723:51.35 Parse Work 13
31863 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 723:17.89 Parse Work 14
31852 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 723:02.83 Parse Work 3
31854 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 722:38.98 Parse Work 5
31856 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 722:40.62 Parse Work 7
31864 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 721:45.72 Parse Work 15
31865 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 725:29.23 Parse Work 16
31838 root      20   0   45.8g   4.6g 545220 S  0.7  3.6 917:20.65 CaptureWork
31839 root      20   0   45.8g   4.6g 545220 S  0.7  3.6 526:43.29 Assem Work
31840 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  28:04.24 Packet Export W
31841 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  42:24.57 Collection ID W
31843 root      20   0   45.8g   4.6g 545220 D  0.7  3.6 303:38.71 Session Write W
31844 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  41:22.46 Session Id Work
31845 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  41:23.96 Session Correla
31849 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  41:19.76 Session Export
32932 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  49:56.33 Request Handler
31771 root      20   0   45.8g   4.6g 545220 S  0.3  3.6  31:57.08 stat updates
```

iotop tool can show Disk i/o activity: rpm is available here and can be installed: `http://mirror.centos.org/centos/7/os/x86_64/Packages/iotop-0.6-4.el7.noarch.rpm`

Example iotop results where decoder session write thread is blocked on IO 99% of decoder IO time and its throughput is just ~30KB/s

iotop –o –d 2

```
Total DISK READ :       0.00 B/s | Total DISK WRITE :     463.05 K/s
Actual DISK READ:       0.00 B/s | Actual DISK WRITE:       0.00 B/s
  TID  PRIO  USER     DISK READ  DISK WRITE  SWAPIN     IO>    COMMAND
 3373 be/4 root        0.00 B/s  431.79 K/s  0.00 % 99.99 % NwDecoder [Packet Write Wo]
 3374 be/4 root        0.00 B/s   31.26 K/s  0.00 % 99.99 % NwDecoder [Session Write W]
```

Resolution:

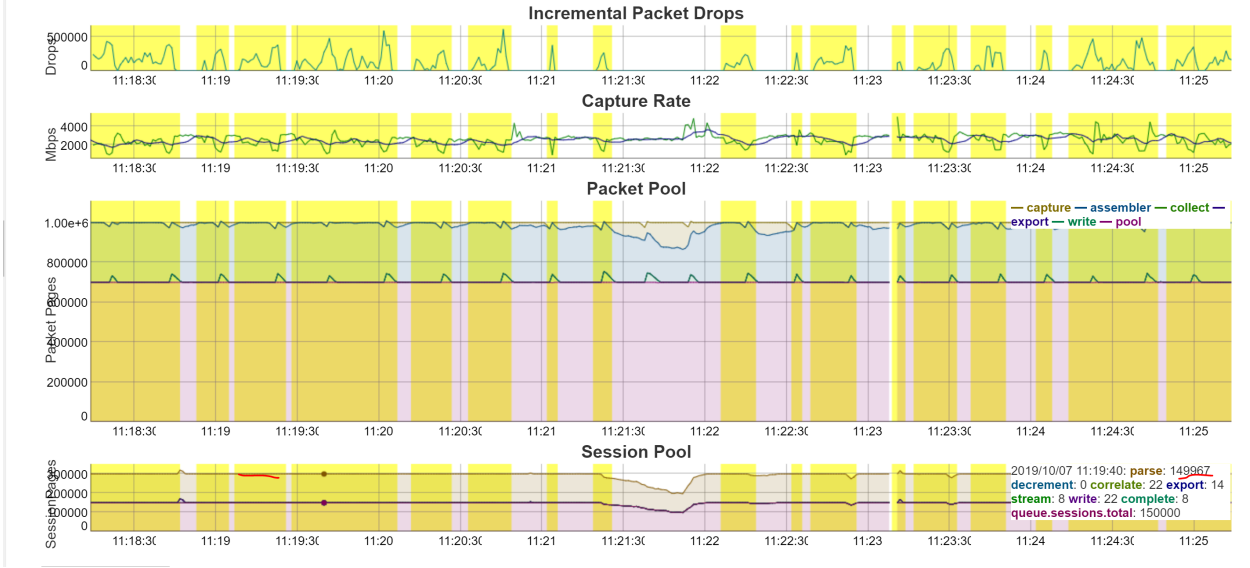
- Make sure database configuration tuning is applied as suggested in section To check and tune the configuration:
- Aggressive aggregators from Decoder
  - Check aggregators using `/decoder whoAgg` or Explore `/decoder | properties | select whoAgg | send`
  - Many aggregators with nice disabled (false) in their configuration can cause session pool backup.
    - For Archiver and Warehouse Connectors nice should be enabled (true).
    - For Concentrator nice should be disabled (false).
  - If there is one device aggregating from Decoder then it would mostly be Concentrator, don't enable nice on Concentrator.

  - If there is an Archiver or Warehouse Connector then set `aggregate.nice=true` On Warehouse Connectors and Archivers aggregating from Decoder.
- Check I/O stats on the Decoder using command "`iostat -mNx 1`". Refer to How do you get statistics on I/O performance?
  - If % `iowait` is > 10% then decoder session db and metadb writes have higher i/o waits
  - If % util for `sessiondb` or `metadb` goes greater than 95 and wMB/s < 1000 , then Disk write throughput is low and the Disk where `sessiondb` and `metadb` exists needs to be replaced.
  - If `iotop` is installed the disk io activity can be monitored through '`iotop -o -d 2`'
  - For a 10G decoder we recommend `sessiondb` and `metadb` write throughput to be 1300 MB/s (~10Gbps) for better session write performance.
- Lot of Content calls to extract meta or content can cause session write issue.
  - Check drops tool logs or `/var/log/messages` or `sosreport logs` for SDK-Content Calls
  - You can also use NwConsole **topQuery** command on messages logs to identify Content calls.
  - Set `/decoder/sdk/config/packet.read.throttle=100` (a higher value) so that packet write would get preference.
  - Check service invoking SDK-Content calls and reduce the content calls.
- Kernel and Driver compatibility issues
  - Check if the firmware is updated according to the Kernel version. If not, update firmware.

Symptom: Session Pool Parsing delays causes packet drops (Parsing issues)

Session Pool Parse - parse (pool.session.parse): If session parse backup increases, then sessions would be queued on assembler and would be causing the drops.

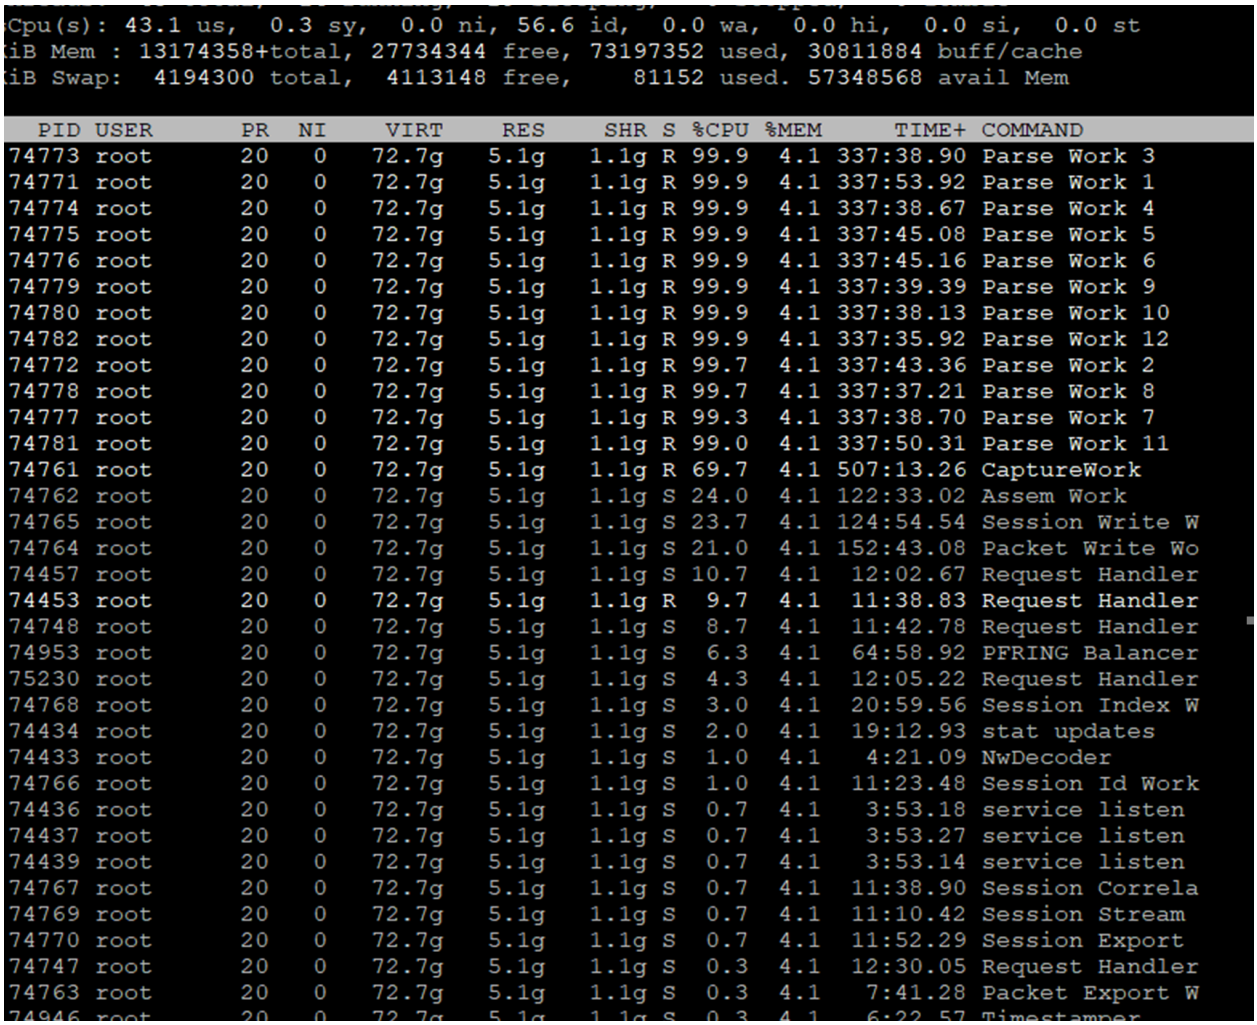
ex: Below screenshot displays there are many sessions waiting to be parsed - Session Pool | parse stat.



Decoder logs would throw warnings like the below:

NwDecoder[42946]: [Packet] [warning] Packet drops encountered, packet assemble (910178/910179): **check session pool (following log), line and session rates**

top -Hp <Decoder pid> would display decoder Parse Work (Parser Threads) are so busy with high CPU usage.



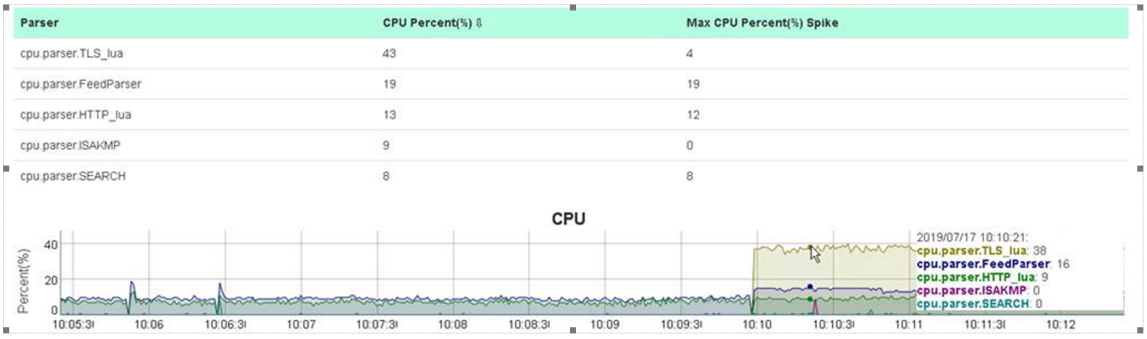
```
Cpu(s): 43.1 us,  0.3 sy,  0.0 ni, 56.6 id,  0.0 wa,  0.0 hi,  0.0 si,  0.0 st
iB Mem : 13174358+total, 27734344 free, 73197352 used, 30811884 buff/cache
iB Swap:  4194300 total,  4113148 free,    81152 used. 57348568 avail Mem

  PID USER      PR  NI    VIRT    RES    SHR S %CPU %MEM     TIME+ COMMAND
74773 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:38.90 Parse Work 3
74771 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:53.92 Parse Work 1
74774 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:38.67 Parse Work 4
74775 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:45.08 Parse Work 5
74776 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:45.16 Parse Work 6
74779 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:39.39 Parse Work 9
74780 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:38.13 Parse Work 10
74782 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:35.92 Parse Work 12
74772 root      20   0   72.7g   5.1g   1.1g R 99.7  4.1 337:43.36 Parse Work 2
74778 root      20   0   72.7g   5.1g   1.1g R 99.7  4.1 337:37.21 Parse Work 8
74777 root      20   0   72.7g   5.1g   1.1g R 99.3  4.1 337:38.70 Parse Work 7
74781 root      20   0   72.7g   5.1g   1.1g R 99.0  4.1 337:50.31 Parse Work 11
74761 root      20   0   72.7g   5.1g   1.1g R 69.7  4.1 507:13.26 CaptureWork
74762 root      20   0   72.7g   5.1g   1.1g S 24.0  4.1 122:33.02 Assem Work
74765 root      20   0   72.7g   5.1g   1.1g S 23.7  4.1 124:54.54 Session Write W
74764 root      20   0   72.7g   5.1g   1.1g S 21.0  4.1 152:43.08 Packet Write Wo
74457 root      20   0   72.7g   5.1g   1.1g S 10.7  4.1  12:02.67 Request Handler
74453 root      20   0   72.7g   5.1g   1.1g R  9.7  4.1  11:38.83 Request Handler
74748 root      20   0   72.7g   5.1g   1.1g S  8.7  4.1  11:42.78 Request Handler
74953 root      20   0   72.7g   5.1g   1.1g S  6.3  4.1  64:58.92 PFRING Balancer
75230 root      20   0   72.7g   5.1g   1.1g S  4.3  4.1  12:05.22 Request Handler
74768 root      20   0   72.7g   5.1g   1.1g S  3.0  4.1  20:59.56 Session Index W
74434 root      20   0   72.7g   5.1g   1.1g S  2.0  4.1  19:12.93 stat updates
74433 root      20   0   72.7g   5.1g   1.1g S  1.0  4.1   4:21.09 NwDecoder
74766 root      20   0   72.7g   5.1g   1.1g S  1.0  4.1  11:23.48 Session Id Work
74436 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1   3:53.18 service listen
74437 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1   3:53.27 service listen
74439 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1   3:53.14 service listen
74767 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1  11:38.90 Session Correla
74769 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1  11:10.42 Session Stream
74770 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1  11:52.29 Session Export
74747 root      20   0   72.7g   5.1g   1.1g S  0.3  4.1  12:30.05 Request Handler
74763 root      20   0   72.7g   5.1g   1.1g S  0.3  4.1   7:41.28 Packet Export W
74946 root      20   0   72.7g   5.1g   1.1g S  0.3  4.1   6:22.57 Timestamper
```

Resolution:

- Enable parser detailed stats on decoder `/decoder/parsers/config/detailed.stats=yes`, wait for new drop instance.
- Click on 

to analyse parsing activity.
- Click on Top Parsers CPU usage, this displays top 5 parsers using higher cpu usages.
  - Try to disable the parsers which are using high CPU usage and wouldn't be parsing the larger percentage of traffic and monitor decoder for drops due to parsing.
    - For ex: parser like TLS_lua , TLD_Lua has higher CPU usage than HTTP_lua (which parsers larger percent of traffic).



Parser	CPU Percent(%) ⇩	Max CPU Percent(%) Spike
cpu.parser.TLS_lua	43	4
cpu.parser.FeedParser	19	19
cpu.parser.HTTP_lua	13	12
cpu.parser.ISAKMP	9	0
cpu.parser.SEARCH	8	8

- Repeat the above step till packet drops are reduced or resolved
- Notify parser Content owner (Customer - if it is their custom content or NetWitness Content team if it is NetWitness parser) so that a fix may be provided for parser issue.

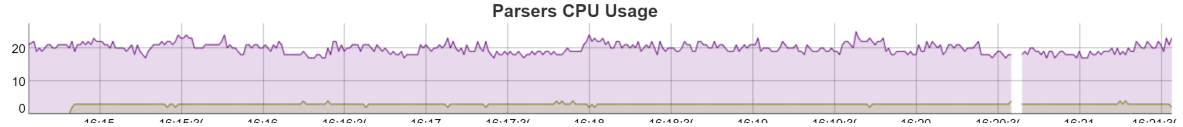
- Click Top Parsers Token Callback counts, this displays top 5 parsers having higher token callback counts during the drop window.
- Try to disable the parsers which are using high Token Callback count and wouldn't be parsing the larger percentage of traffic. Monitor decoder for drops due to parsing.

Top Parsers TokenCallback count

Parser	TokenCallbacks Count ⇩	Max TokenCallbacks Count Spike
TLS_lua	66656555	314517
HTTP_lua	14169879	151699
china_chopper	10648606	116502
rekaf	10604218	115998
Form_Data_lua	8871783	97752

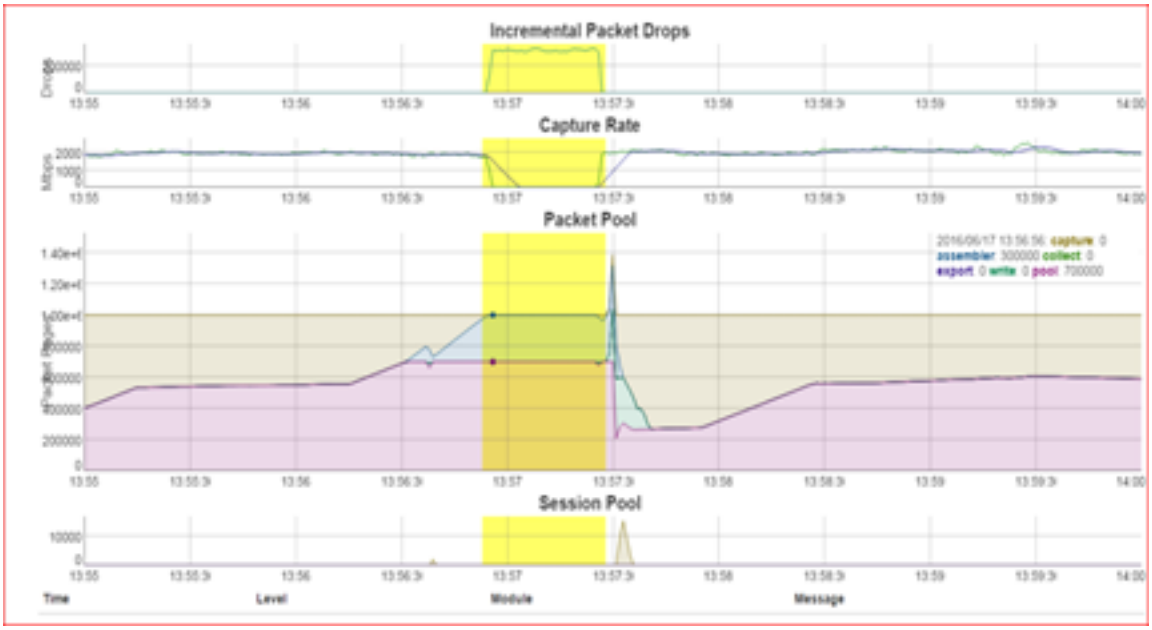
- For ex: parser like TLS_lua has higher Token Callback count than HTTP_lua (which parsers larger percent of traffic)

- Click Top Parsers Meta Callback counts, this displays top 5 parsers having higher meta callback counts during drop window.
- The parsers which displays higher Token Callback count and does not generate much meta and wouldn't parse larger percent of traffic, needs to be disabled.
  - Disable such parsers and monitor Decoder.
- Click Top Parsers Memory Usage, high Lua memory usage (>1 Gb per parser) or memory spike can cause packet drops.
  - Disable Lua Parsers which are using higher memory usage > 1 Gb and monitor Decoder for drops.
- Complex App rules including "regex" or "contains" are CPU intensive
  - Check Parsers CPU Usage graph, if there is high CPU usage for cpu.app.rules stat, then check the app rules configured and tune them.



Symptom: Session Parsing stuck can cause packet drops

Session parsing got stuck at a parser would show symptom of graph like below where queue.sessions.total =1 on Session pool.



Decoder would log warning message like:

NwDecoder[42946]: [Packet] [warning] Packet drops encountered, packet assemble (910178/910179): **check session pool (following log), line and session rates**.

Resolution:

- Enable parser detailed stats on decoder `/decoder/parsers/config/detailed.stats=yes`, wait for new drop instance.
- Click on Analyse Parser Stats to analyse parsing activity.
- Click on Top Parsers CPU usage, this displays top 5 parsers using higher cpu usages.
  - The parser which shows CPU usage when Session Pool `queue.sessions.total =1` would be the one got stuck or looping.
  - Disable the parser and monitor for drops . Notify parser Content owner to get the fix.

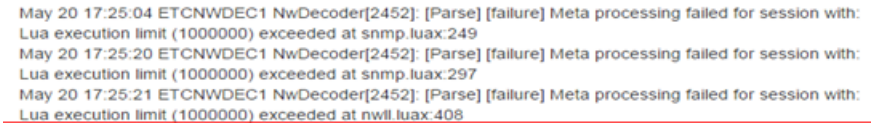
Symptom: Lua Parser Warnings can cause packet drop issues

A lua parser which throws warnings and errors can affect the parse thread and its cpu cycles, hence if there is any lua parser which is throwing warnings then it needs to be fixed.

Resolution:

- Identify and disable the lua parser throwing errors or warnings and monitor decoder.
- Create a content ticket so that content owners can provide a fix or workaround.

ex: Lua Parser throwing parse failure for execution limits exceeding.



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, 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:

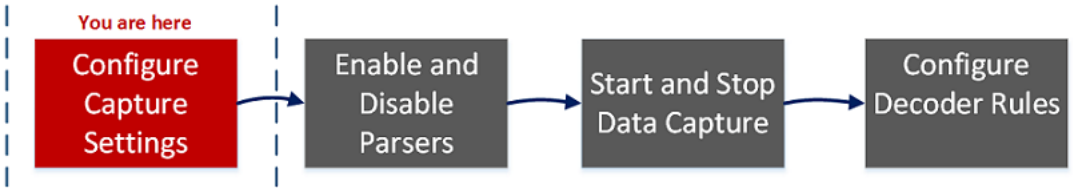
- All base and other protocols (Full Capture - All Protocols)
- Only metadata on all base and other protocols (Capture Meta Only - All Protocols)
- Only metadata on all base protocols and drop all other protocols (Capture Meta on Base Protocols, Drop all other protocols)
- All base protocols and only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

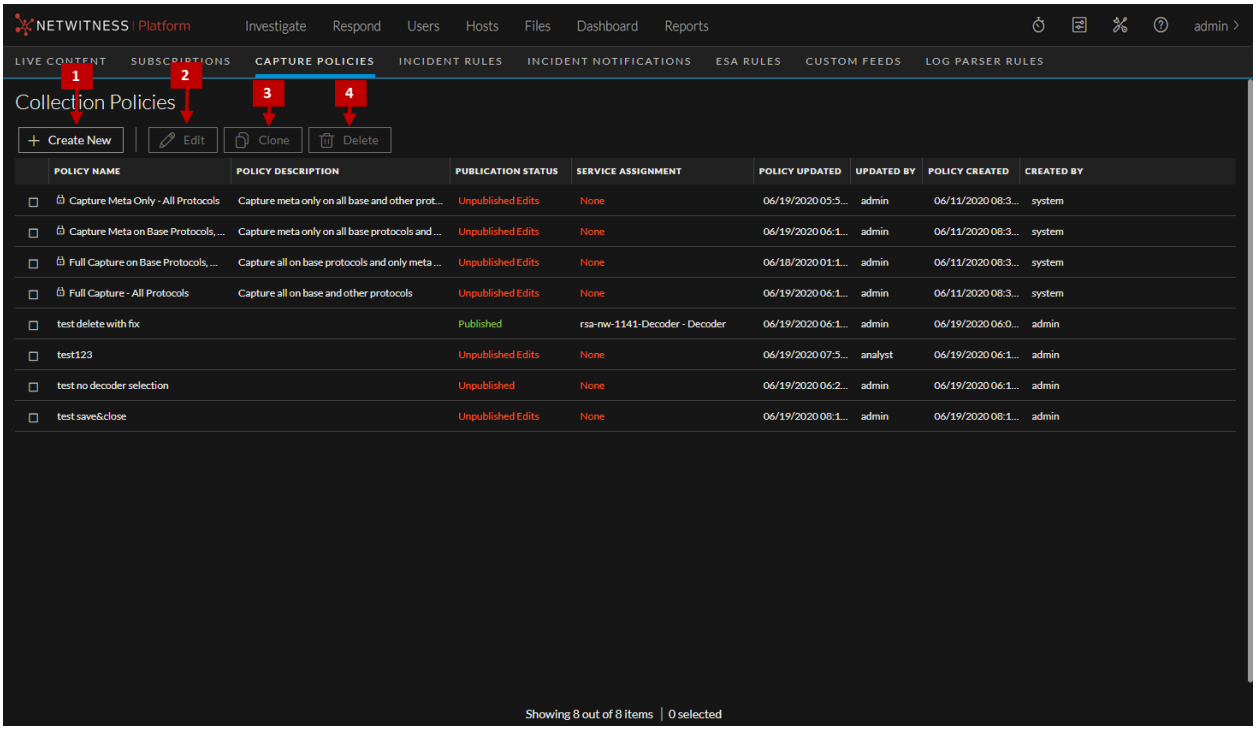
*You can complete these tasks here.

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).
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

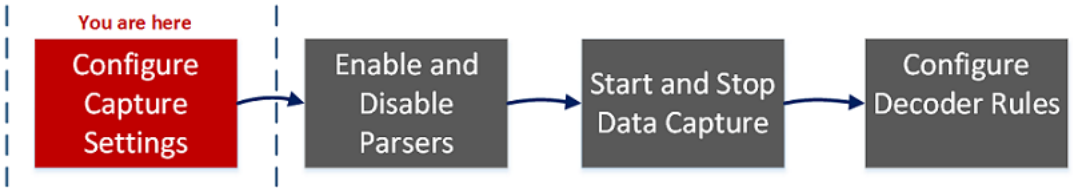
In the Edit Policies view ((Configure) > **Capture Policies** > **Edit**, administrators can configure selective network data collection policies. The panels for creating a new policy and editing an existing policy are the same except that in the editing wizard, the definitions of the existing policy 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 polices 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

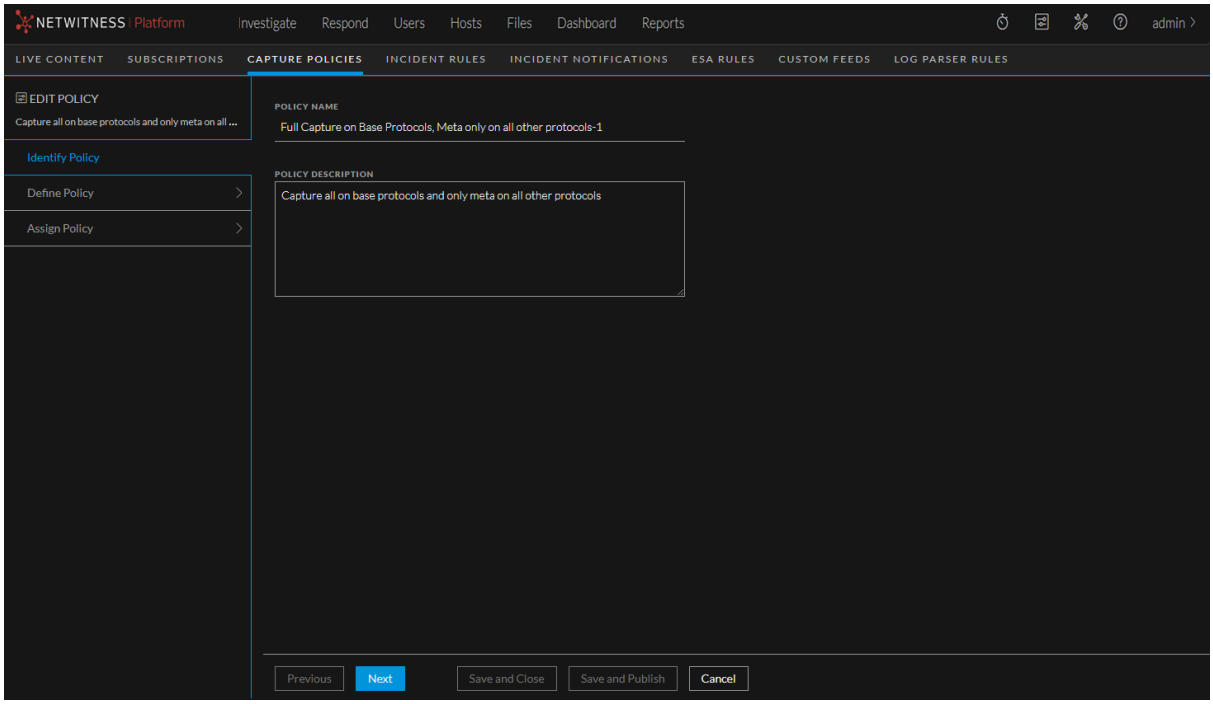
*You can complete these tasks here.

Related Topics

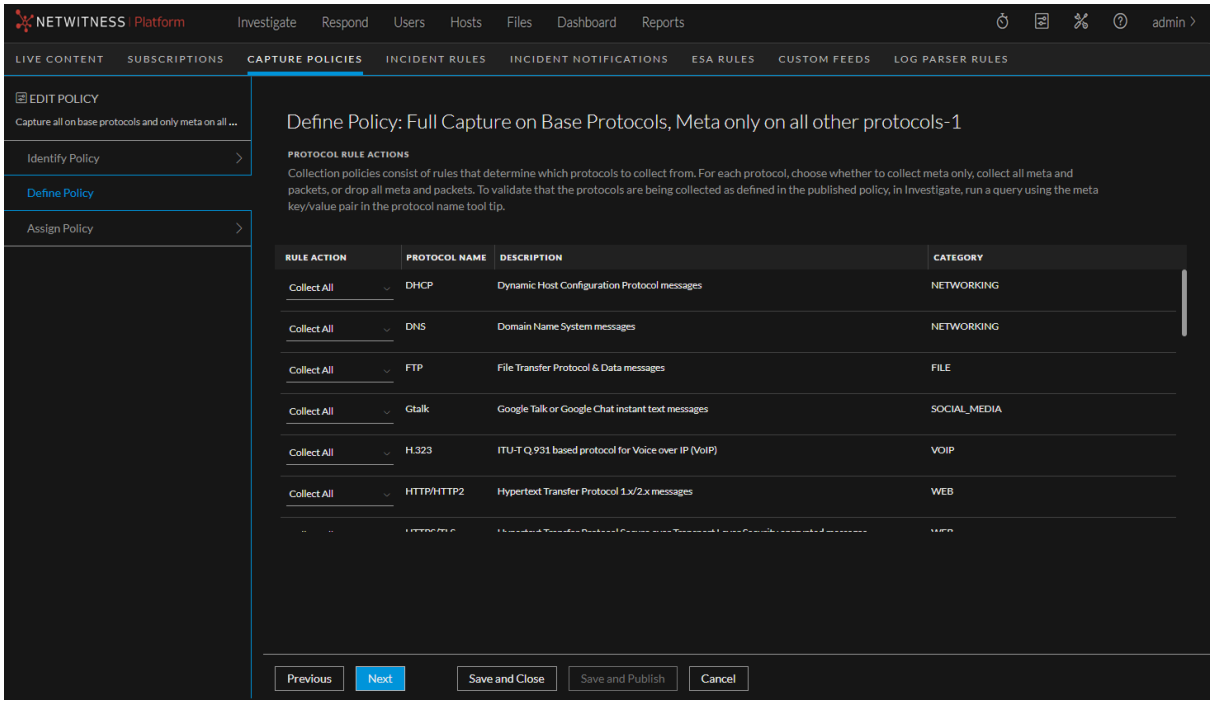
- Services Config View - Capture Policies Tab

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**.



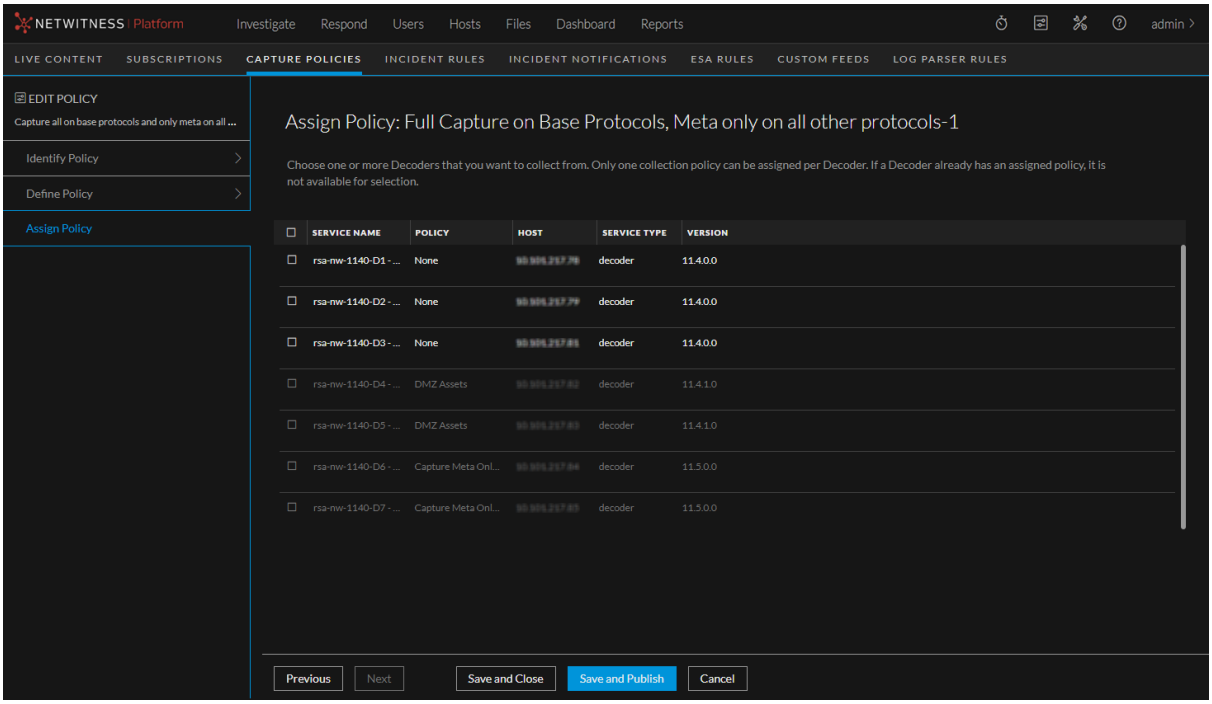
In the Define Policy page, you can define rule actions for the protocols for which you are interested in collecting data. The rule actions are:

- **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**.
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.

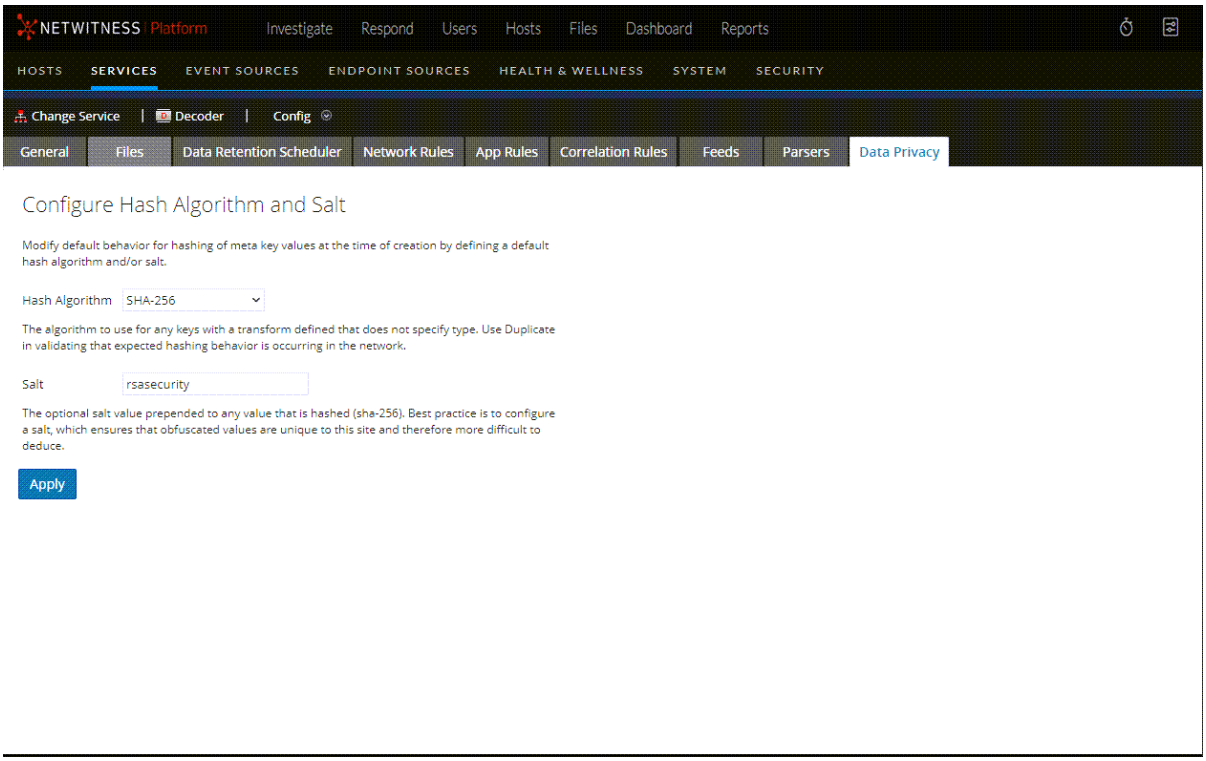
What do you want to do?

User Role	I want to ...	Documentation
Administrator	configure hash algorithm and salt	"Configure the Hash Algorithm and Salt" in the *Data Privacy Management Guide.* (Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder

Quick Look



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 *Data Privacy Management* 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?

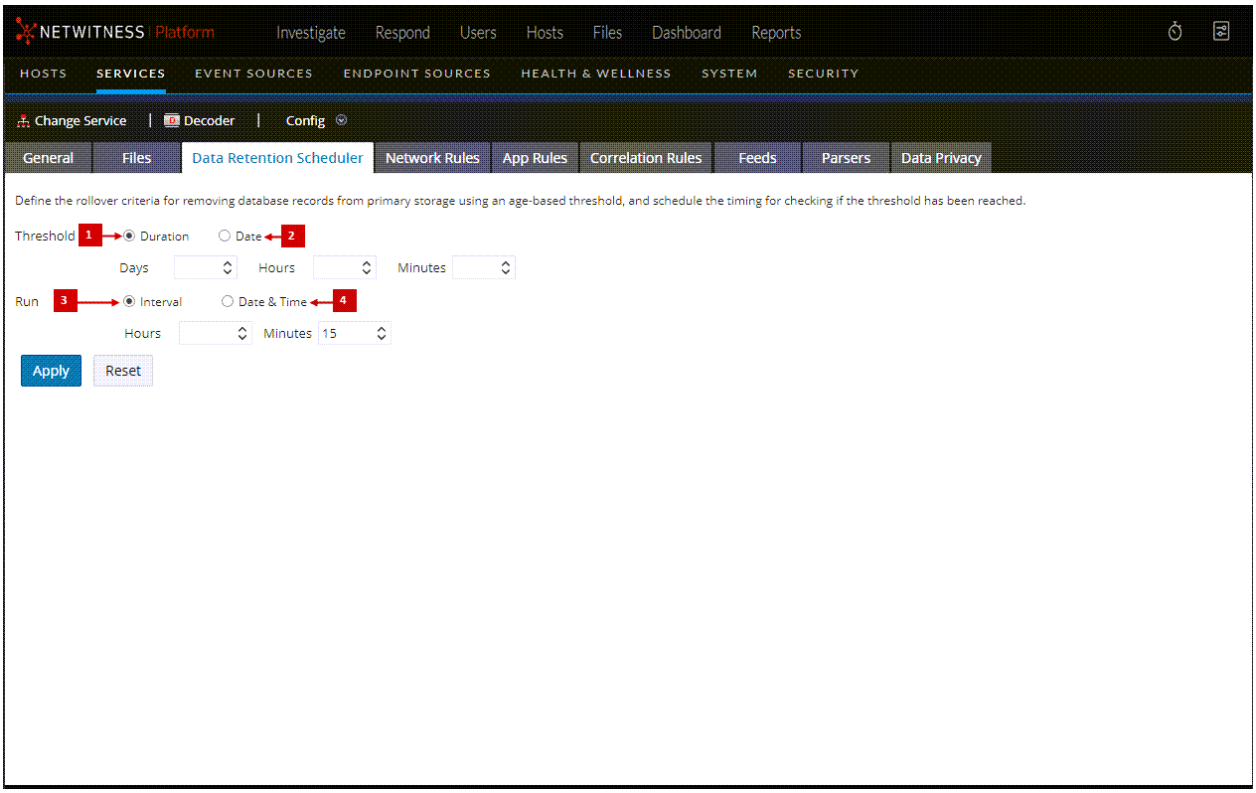
User Role	I want to...	Documentation
Administrator	Schedule the timing to see if the threshold is reached.	Configure Transaction Handling on a Decoder

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.



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.

NetWitness 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?

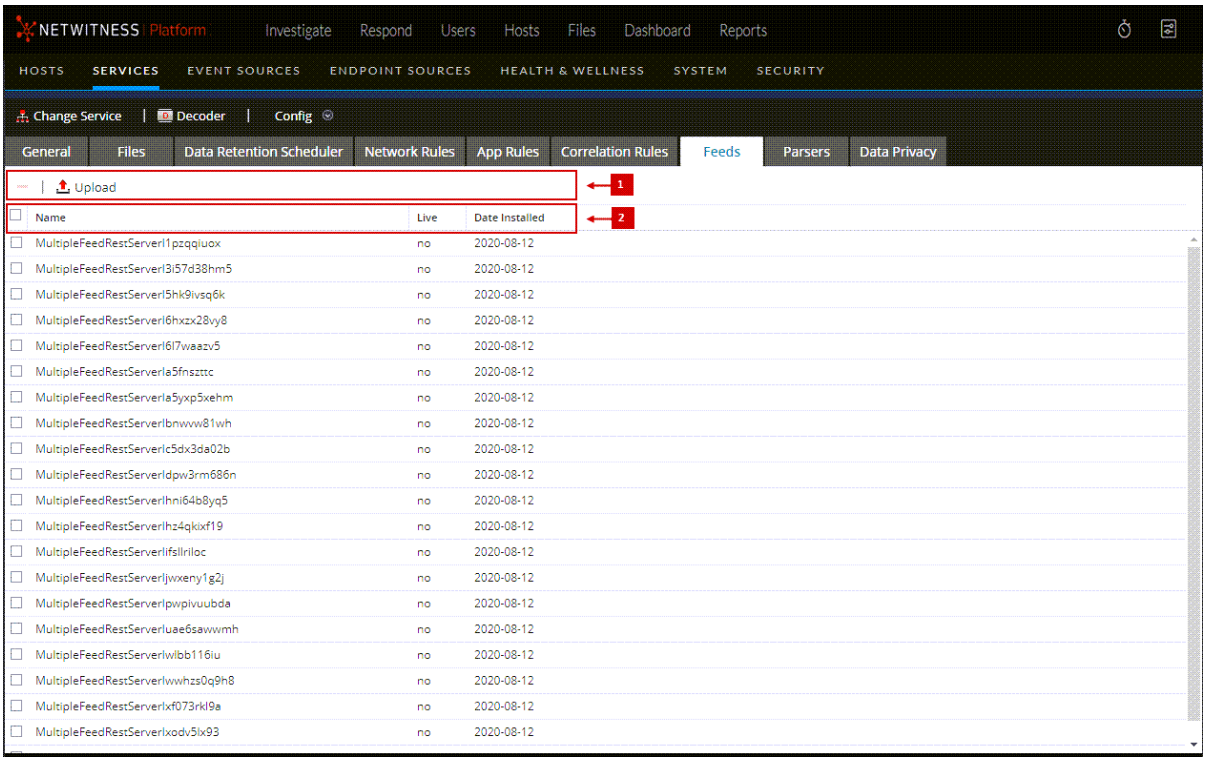
User Role	I want to...	Documentation
Administrator	configure feeds	Configure Parsers and Feeds
Administrator	enable and disable parsers	Enable and Disable Parsers and Log Parsers

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.
Live	Indicates if the feed originated from Live. Possible values are **Yes**, **No**, or **N/A**. • **Yes** = Installed through Live • **No** = Installed through NetWitness • **N/A** = The feed has no attributes file created by NetWitness to track the installation date. The feed may have been installed manually, not through NetWitness or Live Services. Manually installed feeds still function properly.

Column	Description
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?

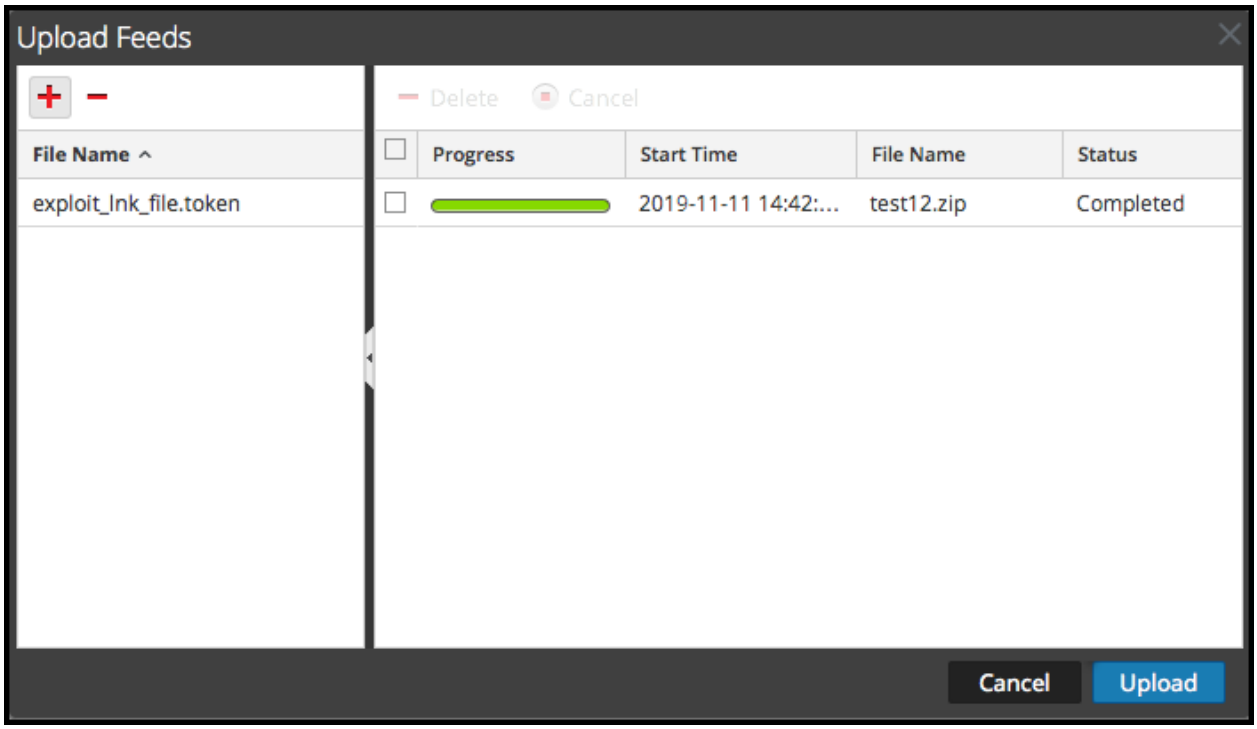
User Role	I want to...	Documentation
Administrator	prepare a list of feeds for upload	Edit, Upload, or Remove a Feed
Administrator	view and delete upload jobs	Edit, Upload, or Remove a Feed

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

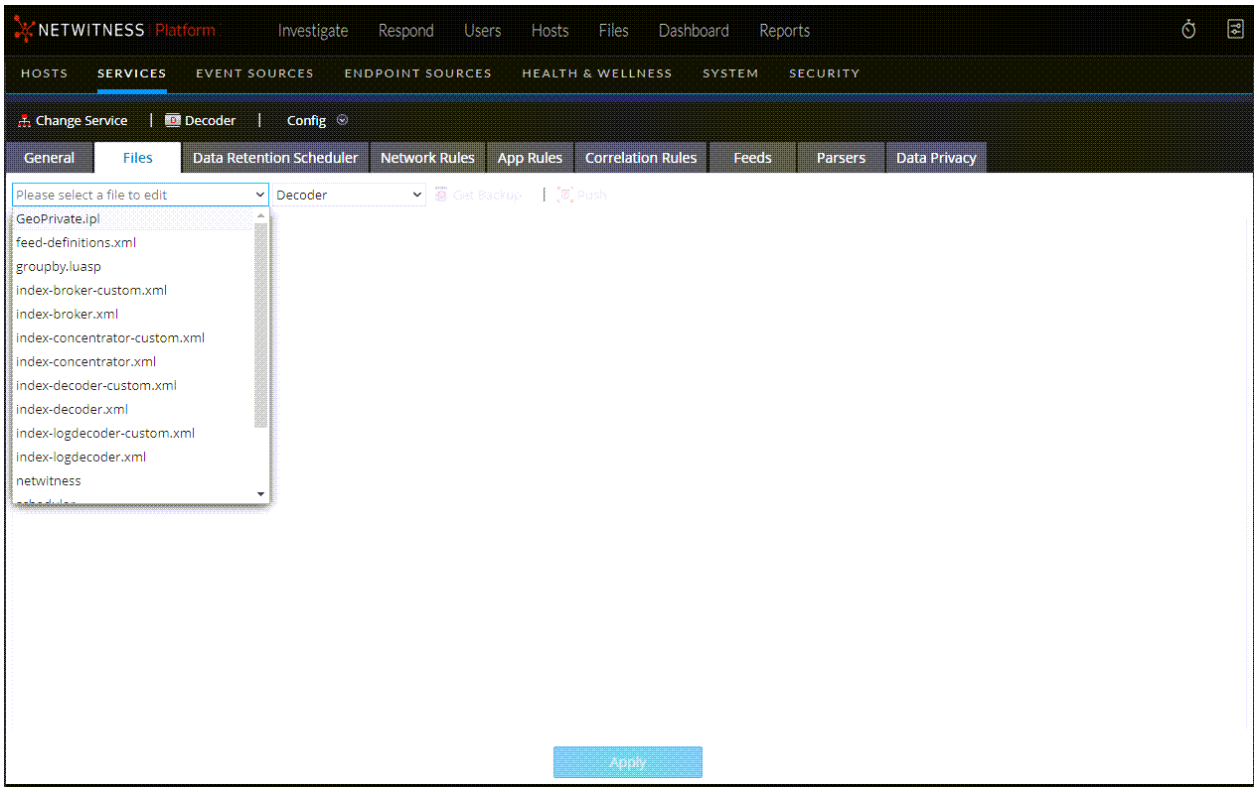
Feature	Description
Cancel	Closes the Upload Feed dialog.
Upload	Starts uploading the feed files listed in the File list. Each feed is listed in a separate row in the Upload job list.

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 NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

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 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

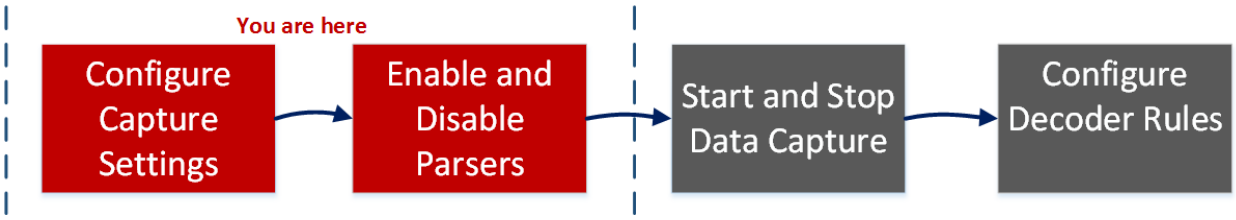
Filename	Description
`GeoPrivate.ipl`	This fixed parser takes the IP addresses and converts them to geographical locations. The locations are displayed through the Google Earth display.
`feed-definitions.xml`	Used to create custom feeds, this is the XML schema used by the Decoder to define a **feed** 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 Content for NetWitness Platform.
`search.xml`	This is the Search Parser configuration file. The Search Parser is a custom parser, used to generate metadata by scanning for pre-defined keywords.
`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

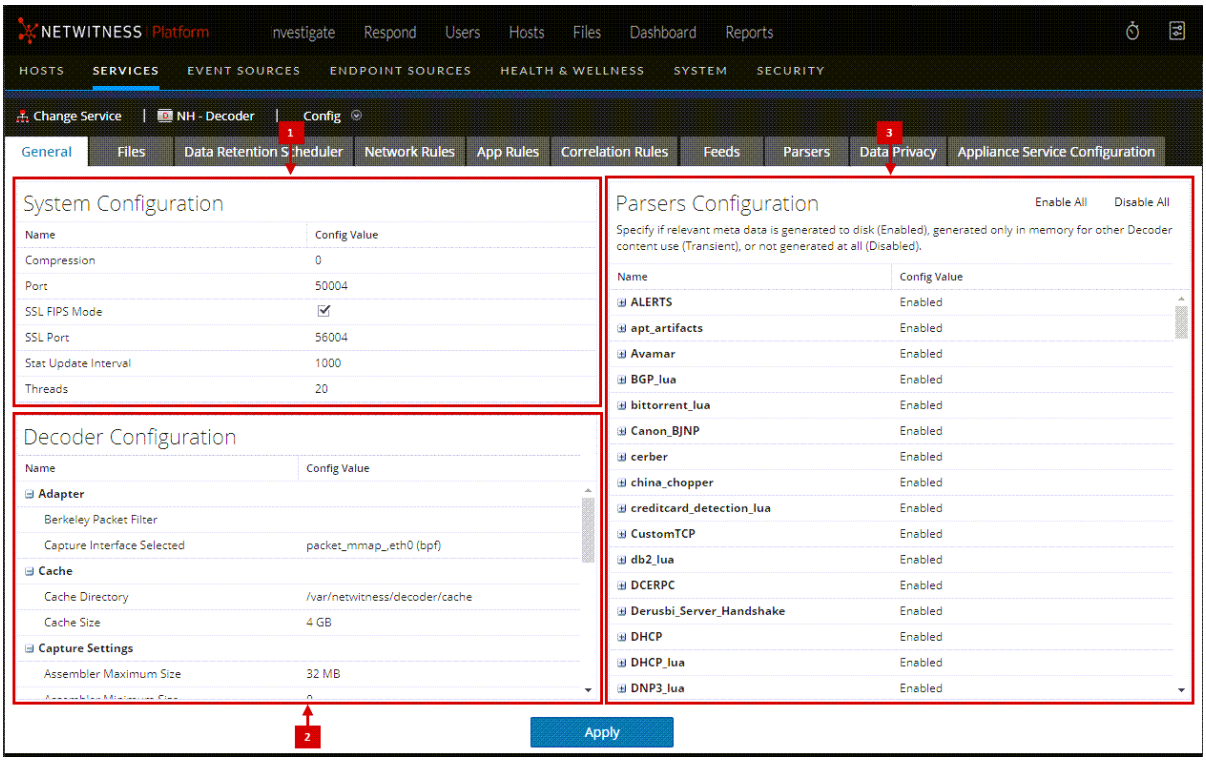
*You can complete these tasks here.

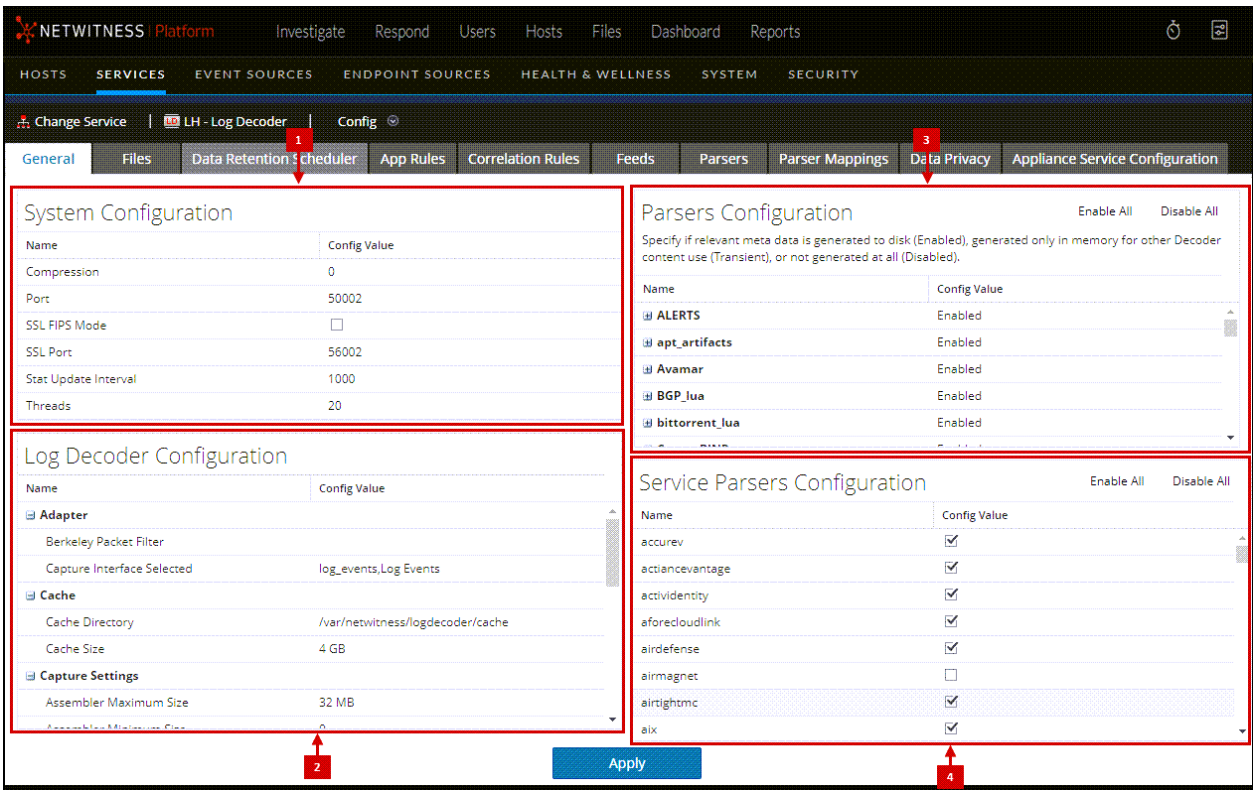
Related Topics

- 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.





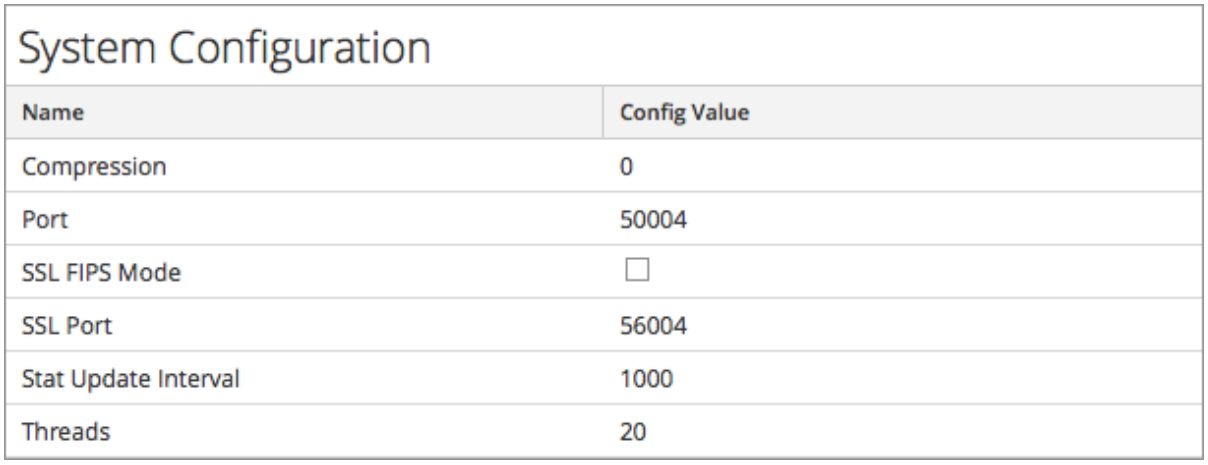
1 System Configuration - Manages service configuration for a Decoder or a Log Decoder.

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

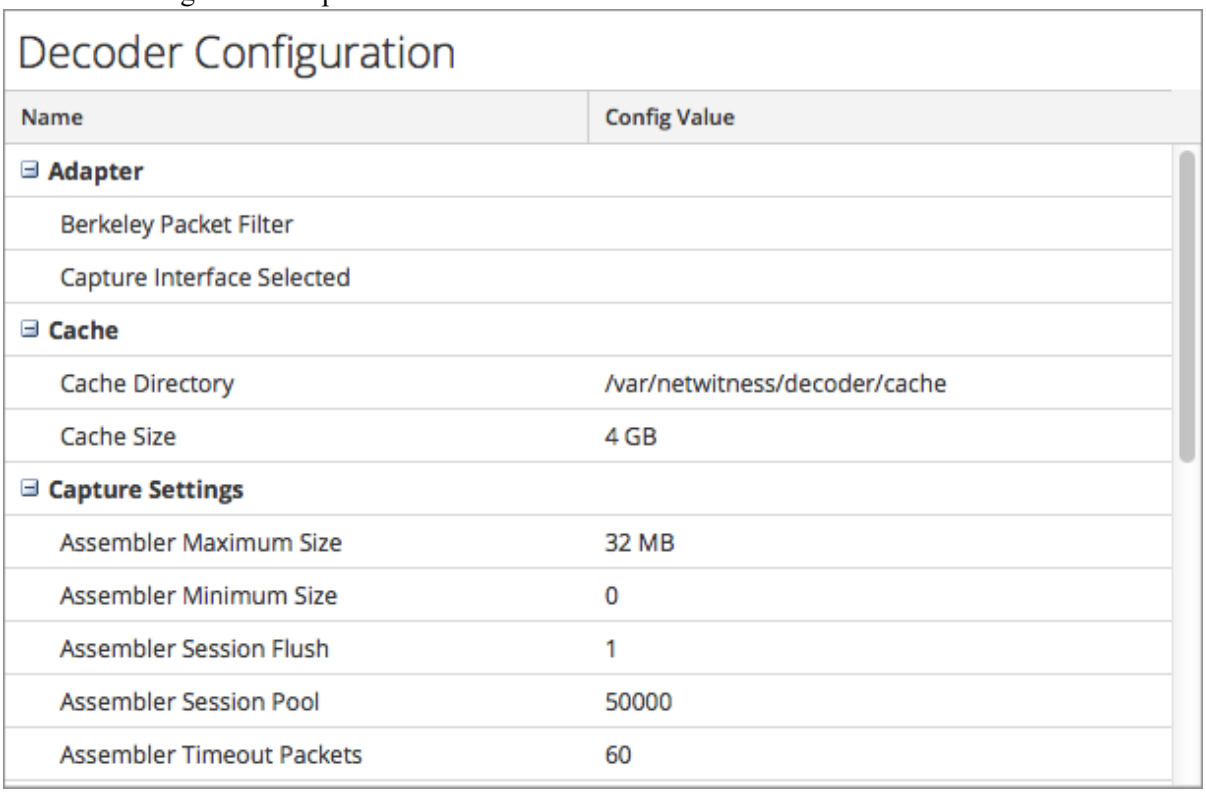
Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☐
SSL Port	56004
Stat Update Interval	1000
Threads	20

The System Configuration section has these parameters.

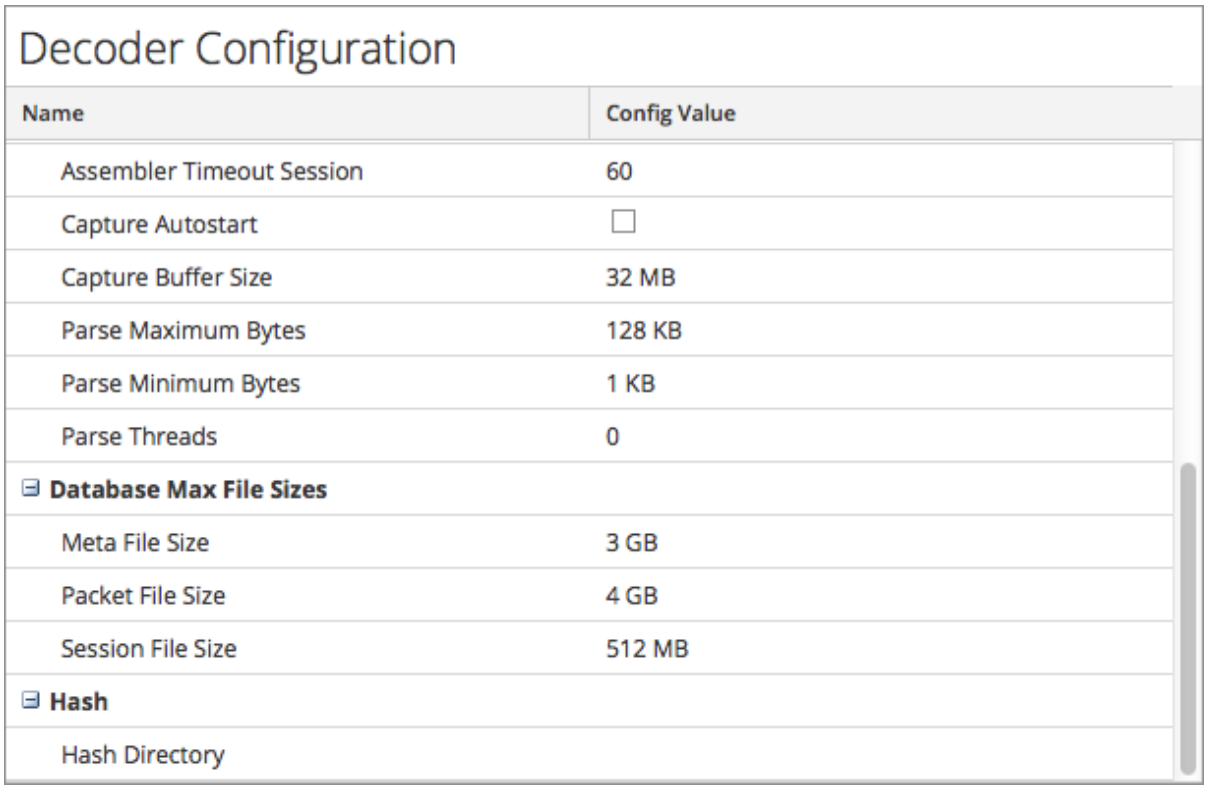
Parameter	Description
Compression	The minimum number of bytes that must be transmitted per response before compression. A setting of 0 disables compression. The default value is **0**. A change in value is effective immediately for all subsequent connections.
Port	Determines the port used by the service. **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.



Scrolling to the bottom of the section reveals these additional Decoder Configuration parameters.



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 `/var/netwitness/decoder/cache`. 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. T 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.

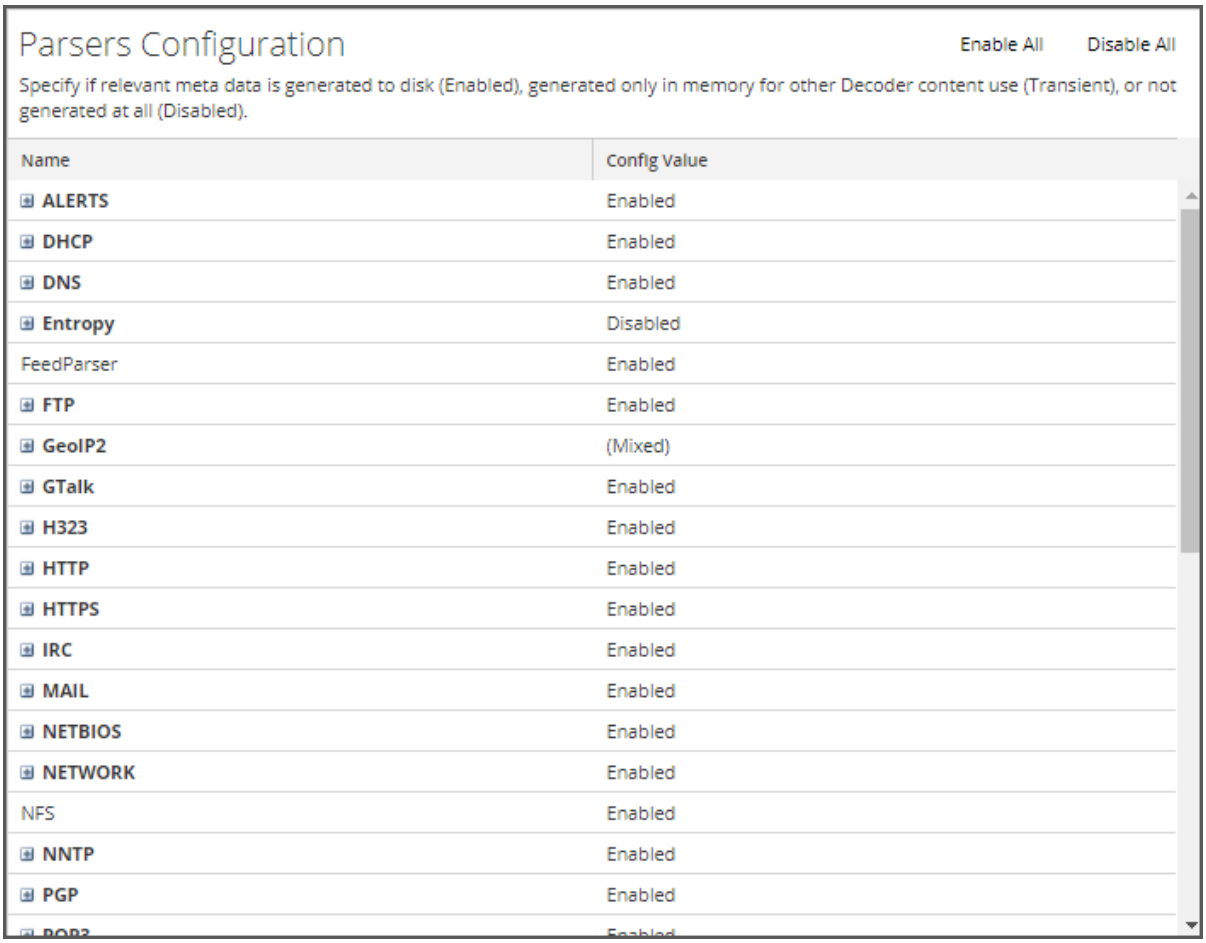
Hash Section

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.

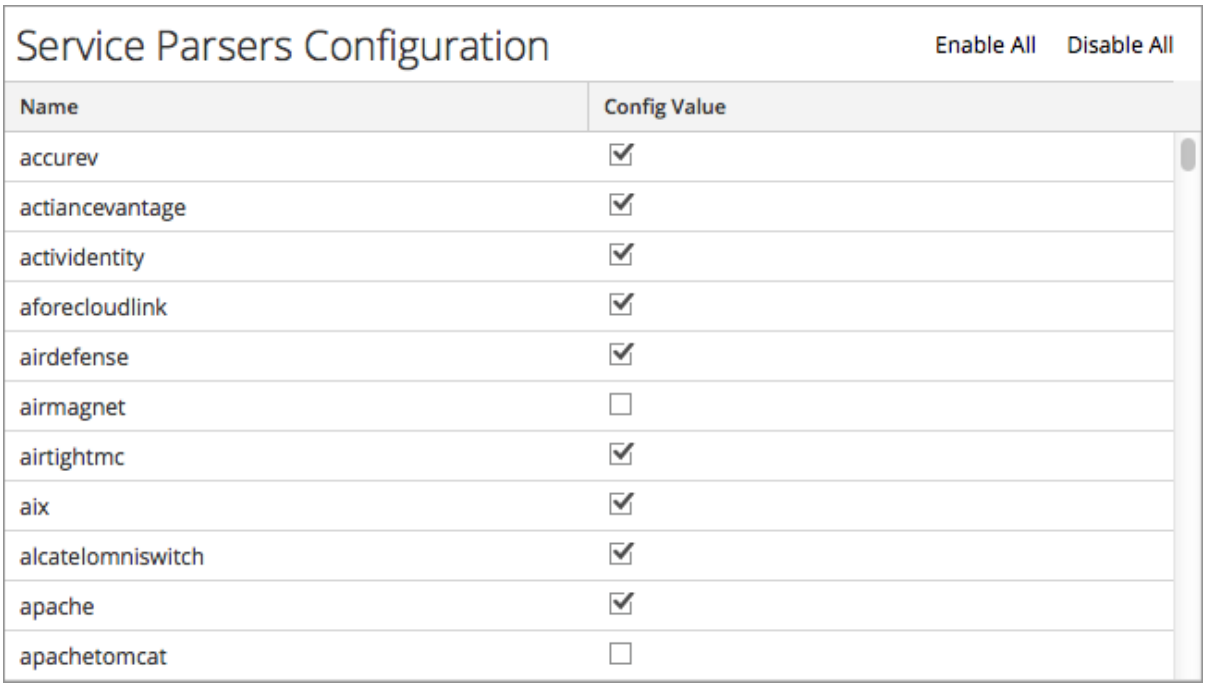
Parsers Configuration Panel

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.



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?

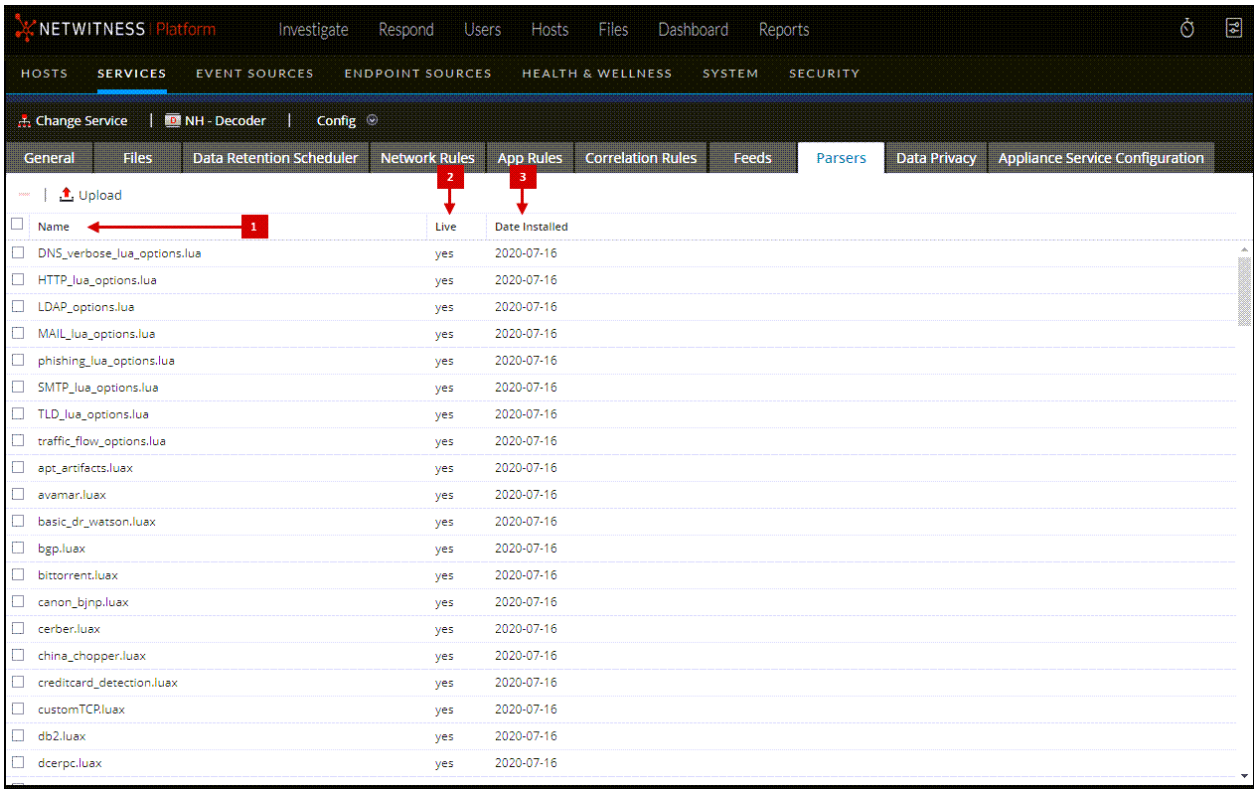
User Role	I want to...	Documentation
Administrator	View deployed parsers.	Enable and Disable Parsers and Log Parsers
Administrator	Upload parsers to a Decoder or Log Decoder.	Enable and Disable Parsers and Log Parsers

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?

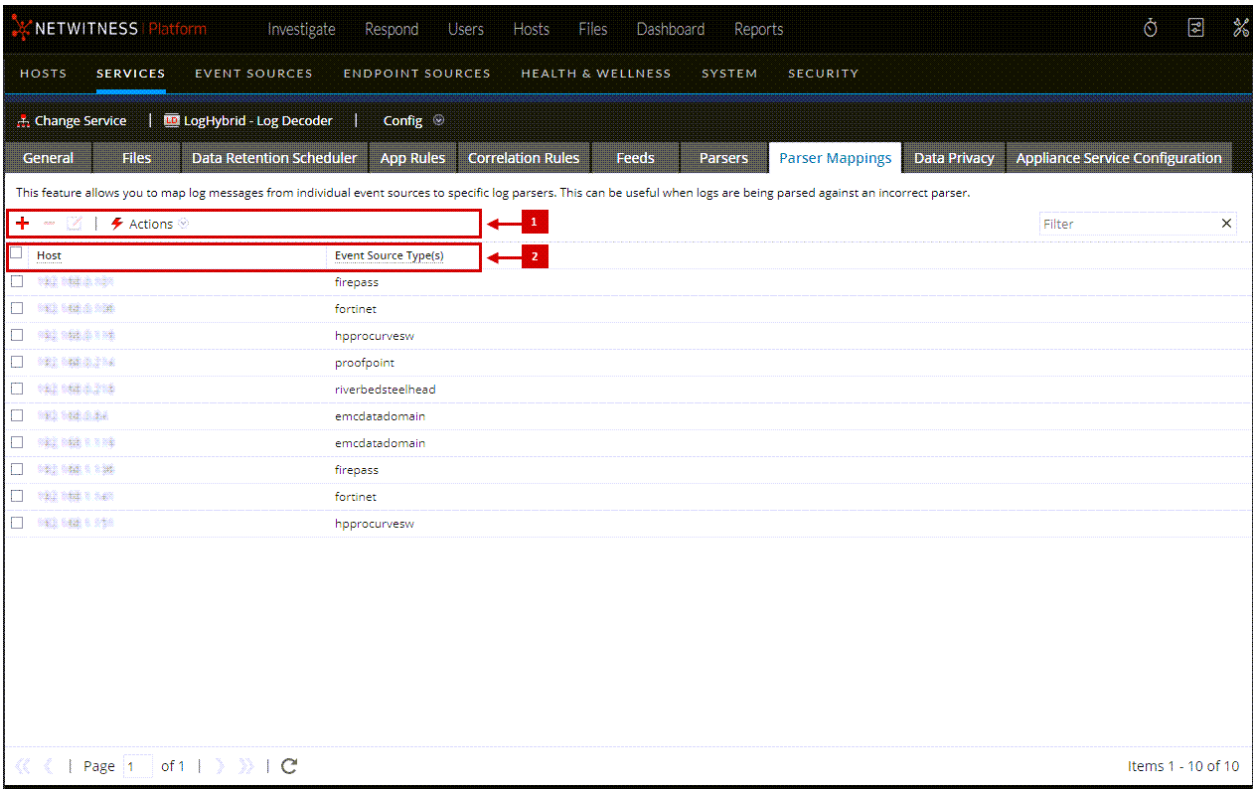
User Role	I want to...	Documentation
Administrator	Manage IPs for Event Source Mapping.	Enable Parser Mappings

Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder

Quick Look

This is an example of the tab.



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?

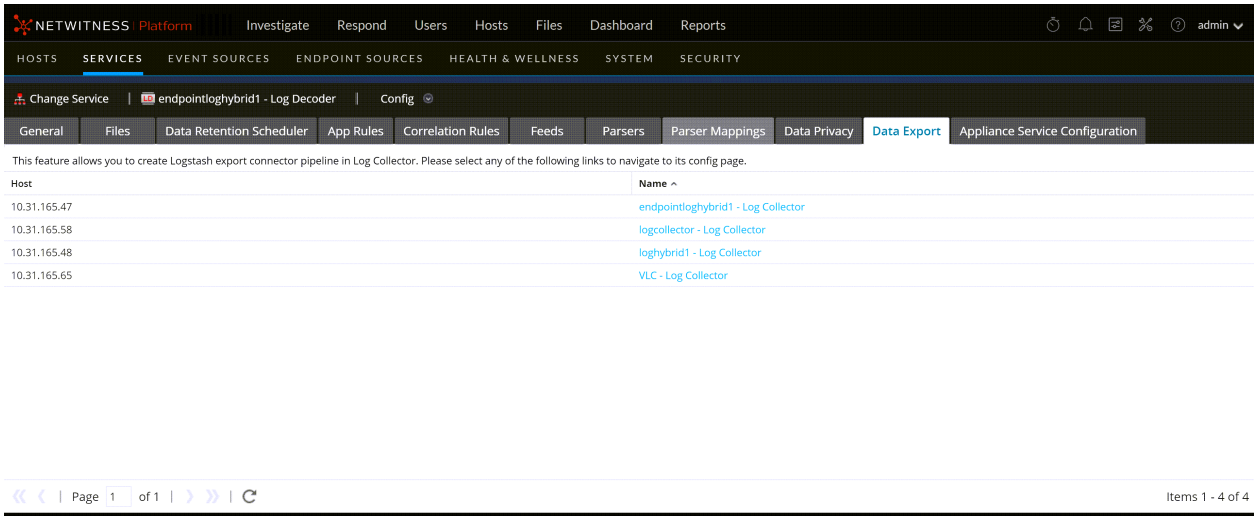
User Role	I want to...	Documentation
Administrator	Configure Export Connector	Configure Logstash Event Sources in NetWitness

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 data export list displays all the Log Collectors in your deployment.

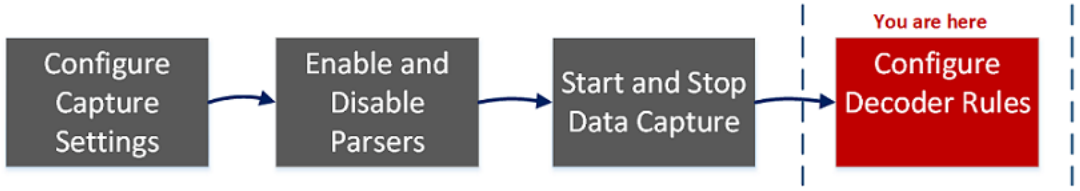
Parameter	Description
Host	Displays the IP address of the host.
Name	Displays the name of the Log Collector.

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

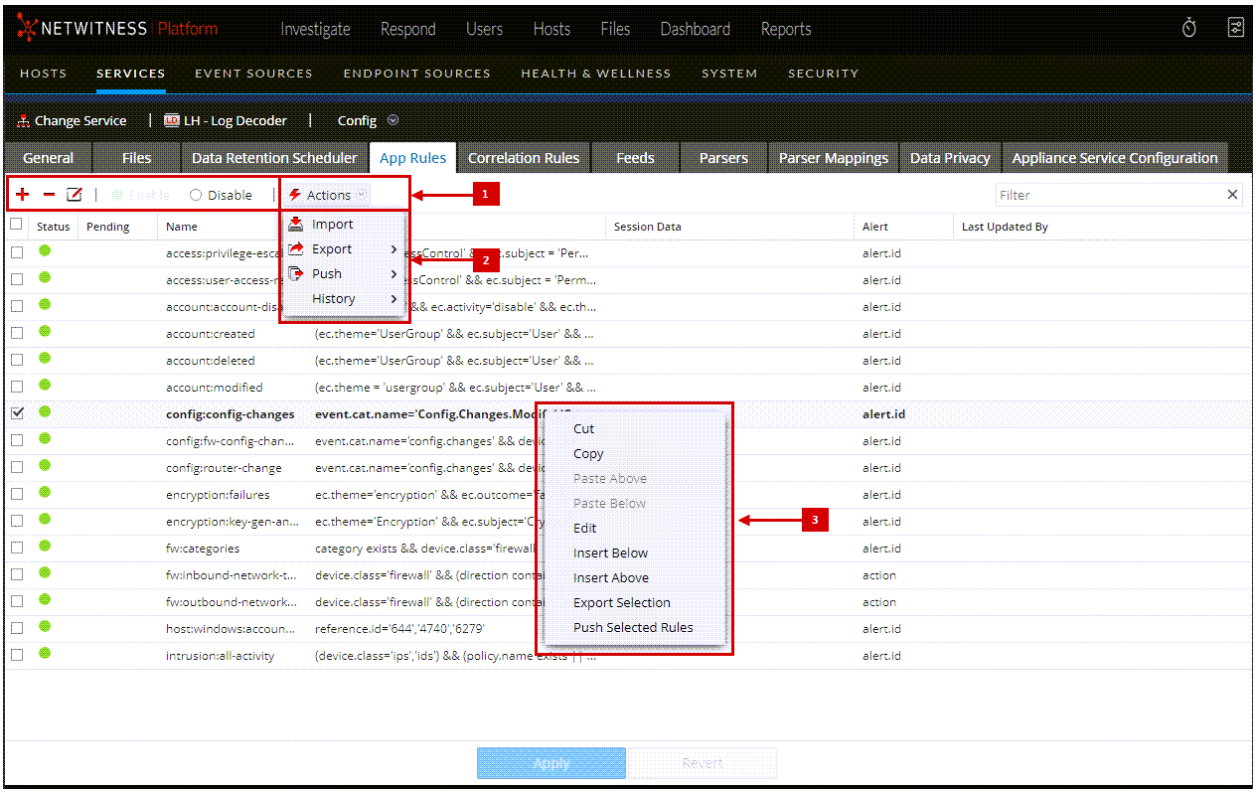
*You can complete these tasks here.

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.

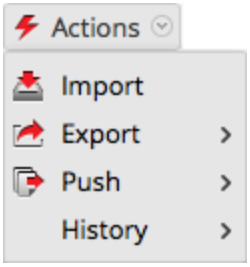


Feature	Description
Actions	Displays the **Actions** menu.
+	Adds a new rule to a service.
−	Deletes a rule from a service.
	Allows rule modification.
Disable	Disables a rule (without deleting the rule).
Enable	Enables (reactivates) a rule.

Feature	Description
Filter	The input field for a search string. NetWitness 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. 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.

Option	Description
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 applies application rules at the session level.

What do you want to do?

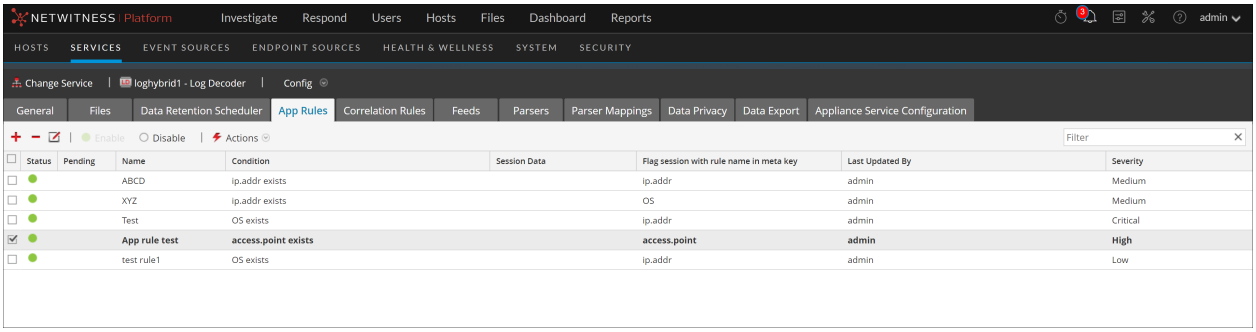
User Role	I want to...	Documentation
Administrator	add or edit application rules	Configure Application Rules

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.

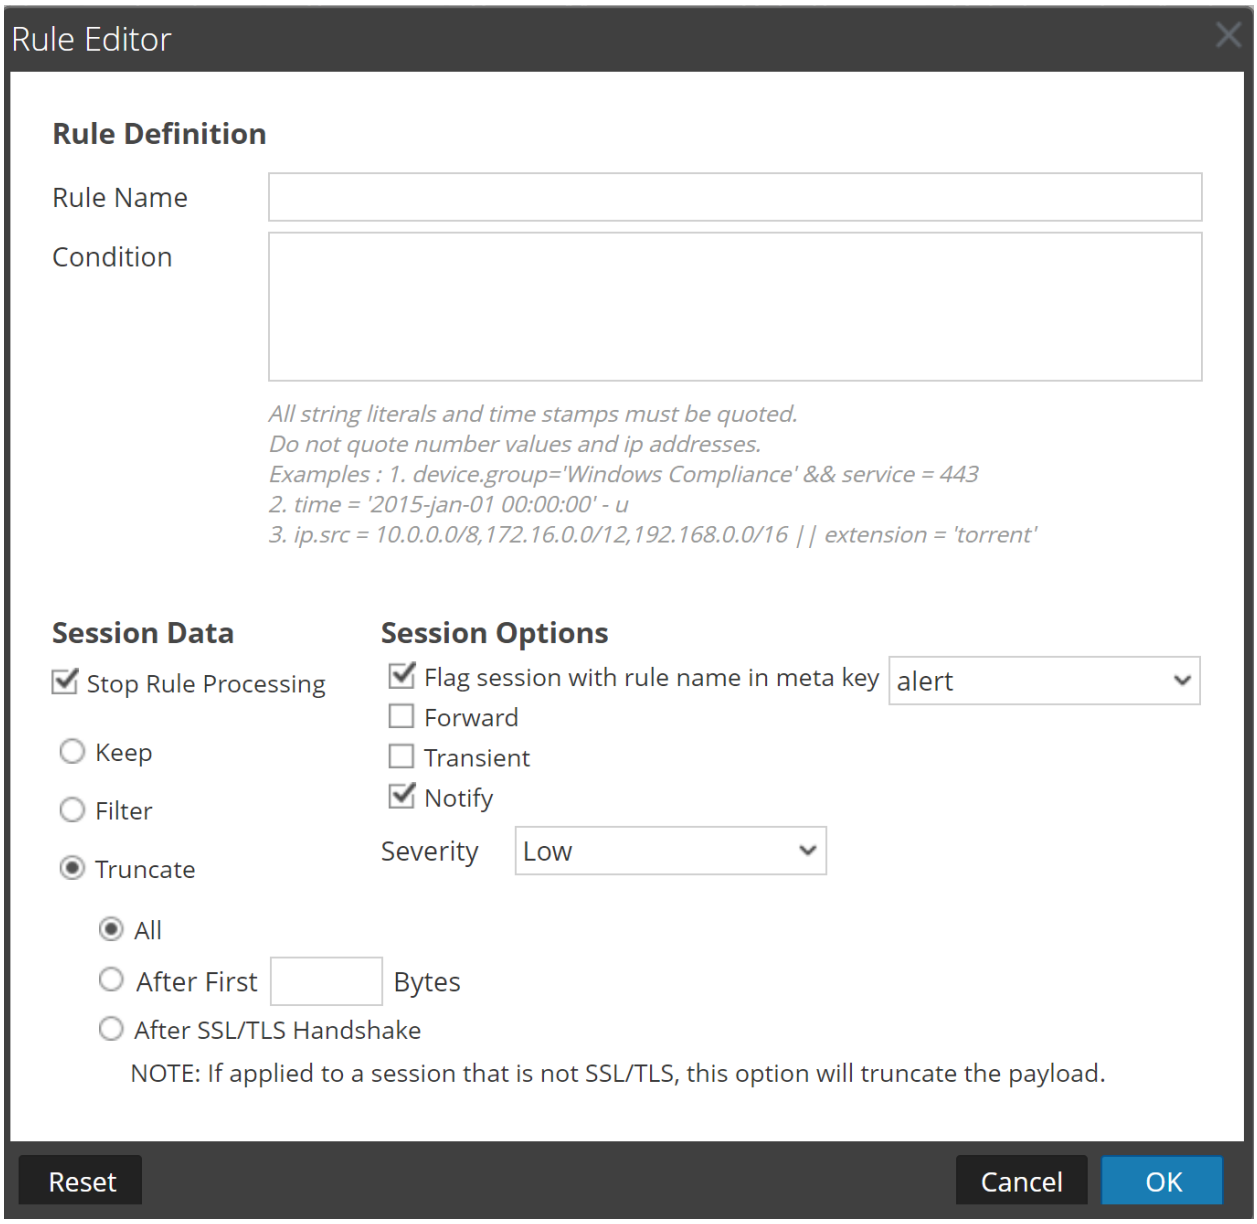


Column	Description
Status	This column indicates whether the rule is enabled or disabled with a circle icon. If the circle is filled green, the rule is enabled. If the circle is empty, the rule is disabled.
Pending	This column indicates whether a rule has pending changes. Rules that are currently active 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.

Column	Description
Session Data	This column displays the Session Data action taken when a packet matches the rule. Possible values are **Filter**, **Keep**, or **Truncate**.
Flag session with rule name in meta key	This column displays the name of the custom alert that the Decoder generates when metadata matches the rule.
Last Updated By	This column displays the user ID of the user who as last updated the rule.
Severity	This column displays the severity of the rule. The available options are **Low**, **Medium**, **High** and **Critical**.

Rule Editor Dialog

The following figure shows the Rule Editor dialog for an application rule.



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.
Flag session with rule name in meta key	If checked, generates a custom alert when metadata matches the rule. You can select the metadata of the alert in the drop-down list.
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.
Notify	Enable the option to select the **Severity** level for the application rule. The available options are **Low**, **Medium**, **High**, and **Critical**.

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 applies correlation rules over a configurable sliding time window.

What do you want to do?

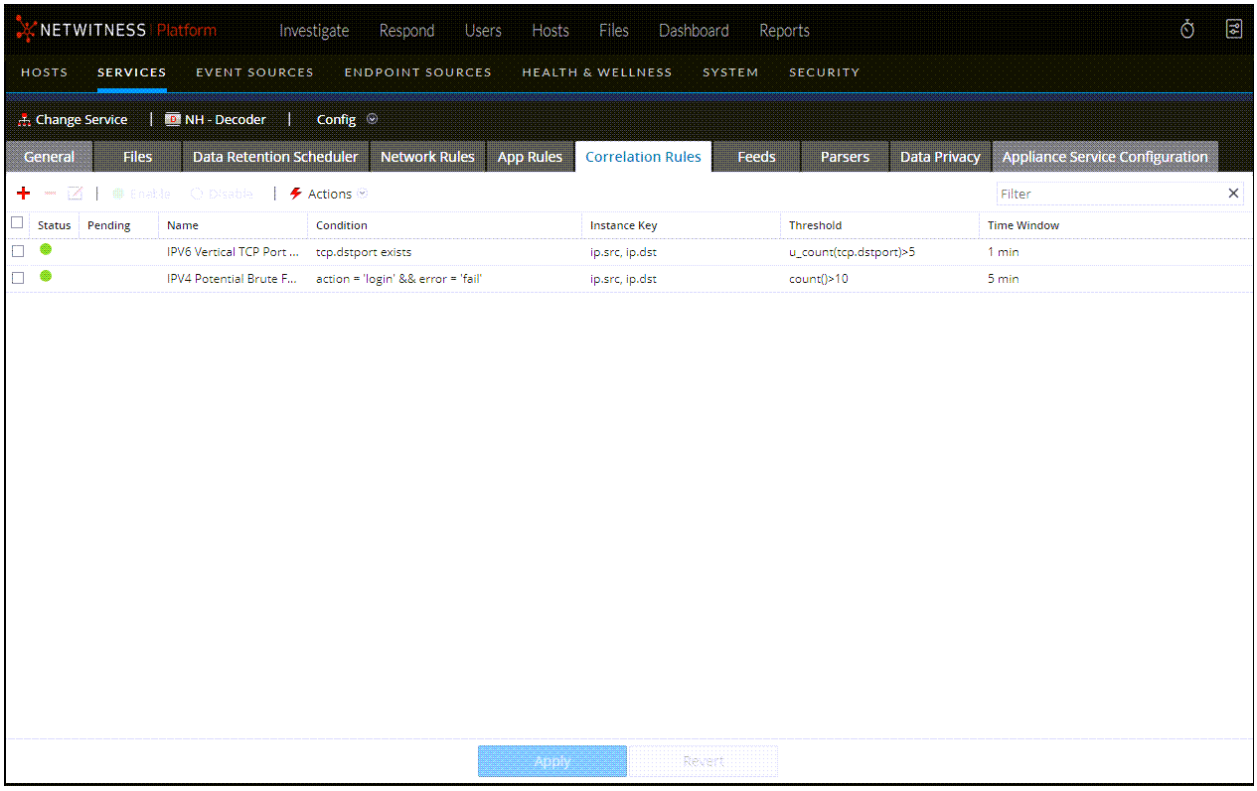
User Role	I want to...	Documentation
Administrator	add or edit a correlation rule	Configure Correlation Rules

Related Topics

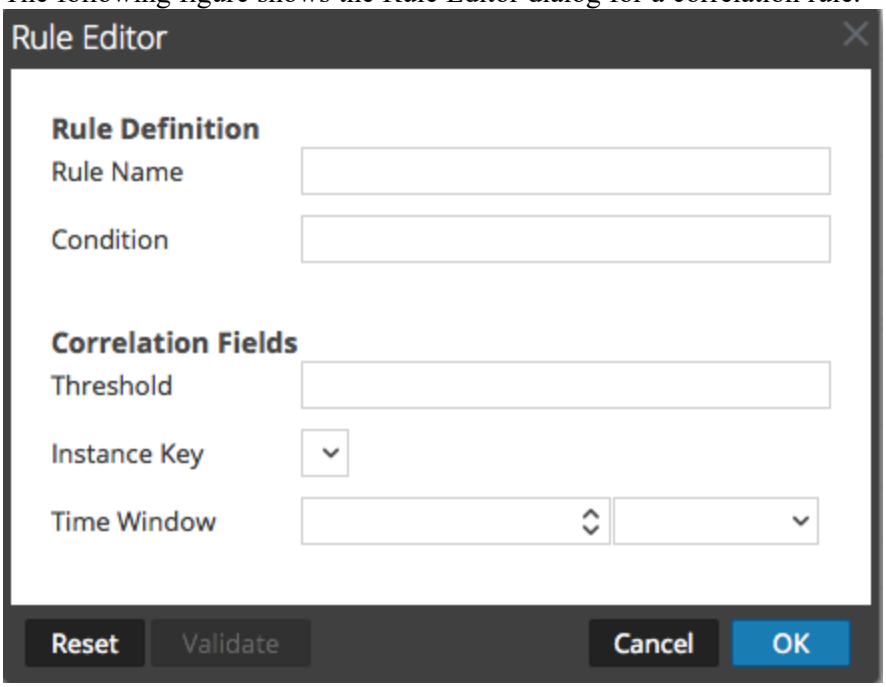
- Configure Common Settings on a Decoder
- Decoder and Log Decoder Quick Setup
- Configure Decoder Rules
- Services Config View - Rules Tabs

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 ip.src or a compound primary key such as `ip.src,ip.dst`.
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: • `u_count(associated_key)` = the count of unique values of the specified key. A key is required. • `sum(associated_key)` = the values of the specified key. a key is required. • `count()` = 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.

Column	Description
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 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?

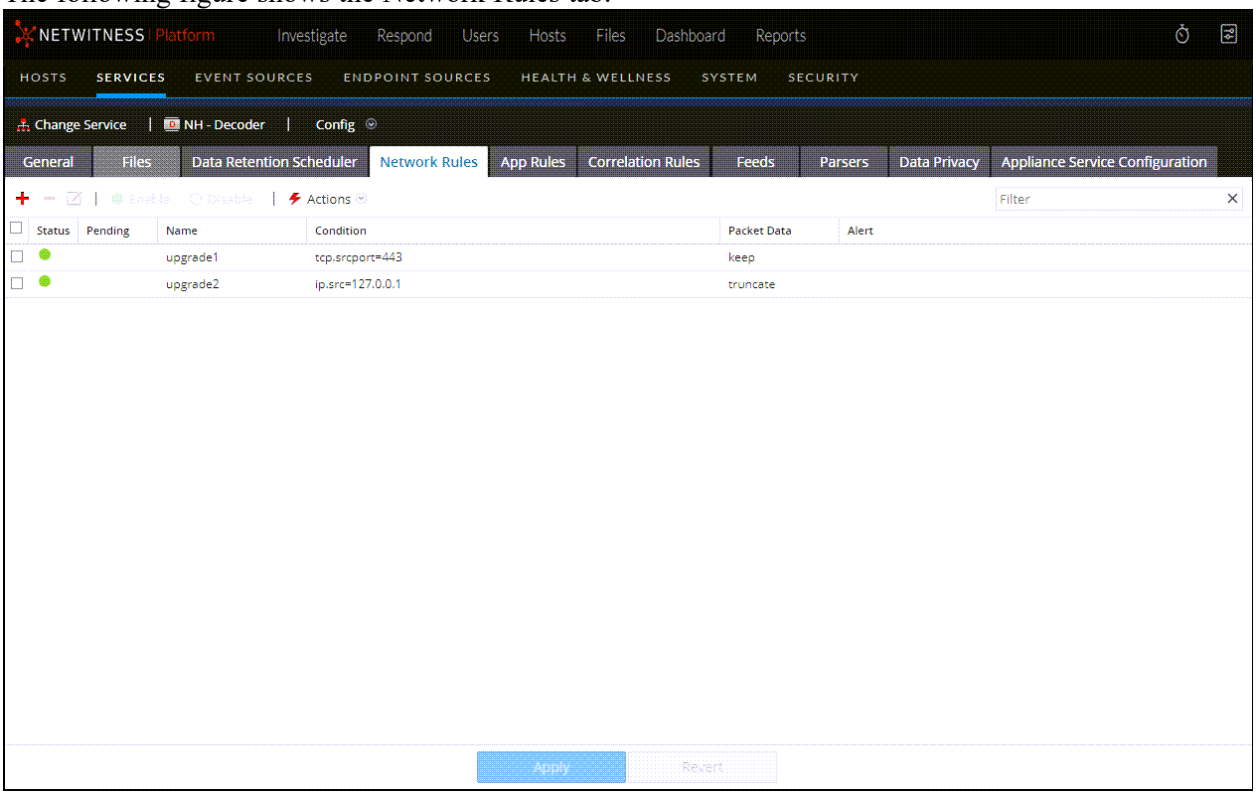
User Role	I want to...	Documentation
Administrator	add, edit, or fix network rules	Configure Network Rules

Related Topics

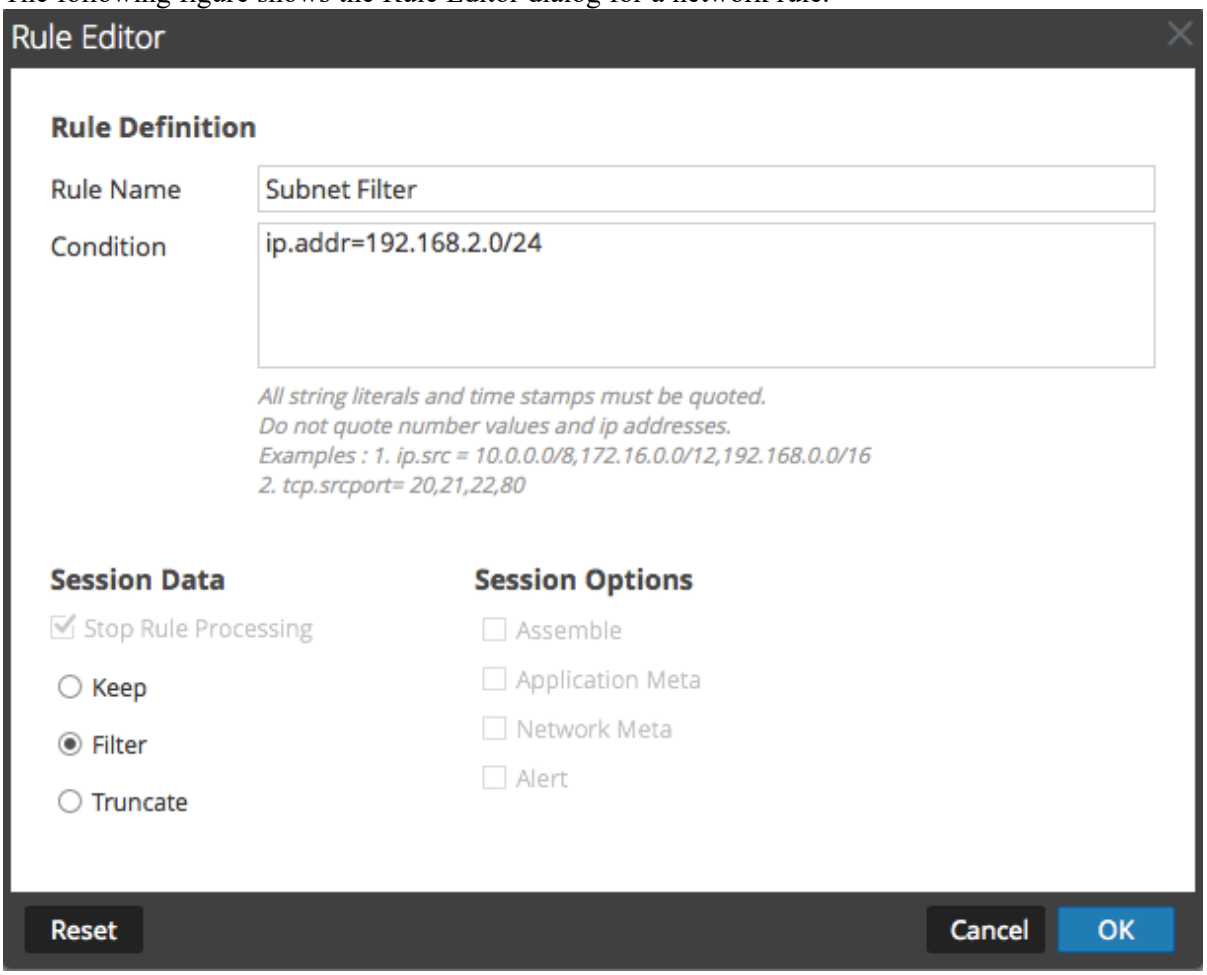
- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder
- Configure Decoder Rules
- Services Config View - Rules Tabs

Quick Look

The following figure shows the Network Rules tab.



The following figure shows the Rule Editor dialog for a network rule.



The following table describes the columns in the Network Rules grid.

Column	Description
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.
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**.

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 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.

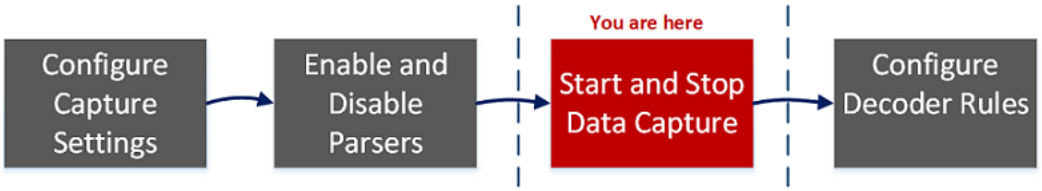
Services System View - Decoders

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*	*Hosts and Services Getting Started Guide*
Administrator	configure rules	Configure Decoder Rules

*You can complete these tasks here.

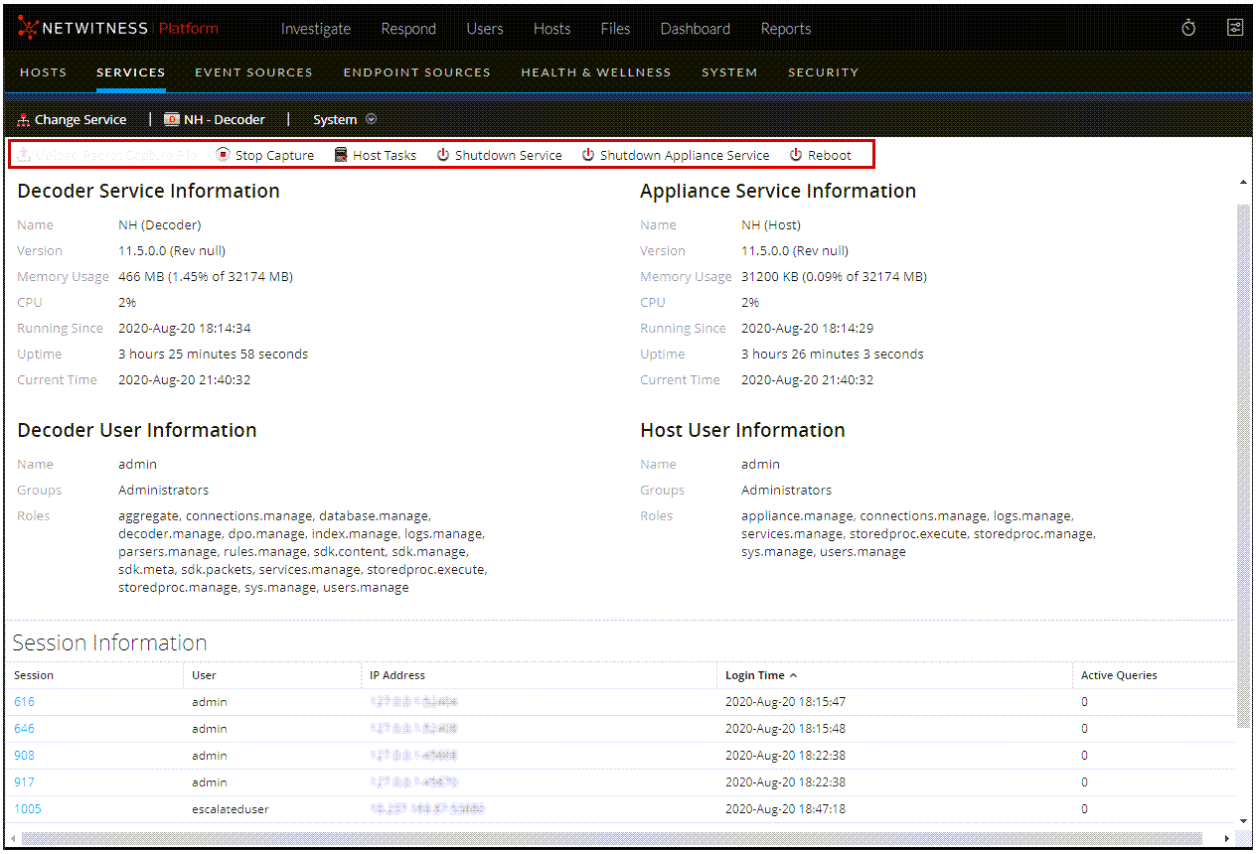
Related Topics

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

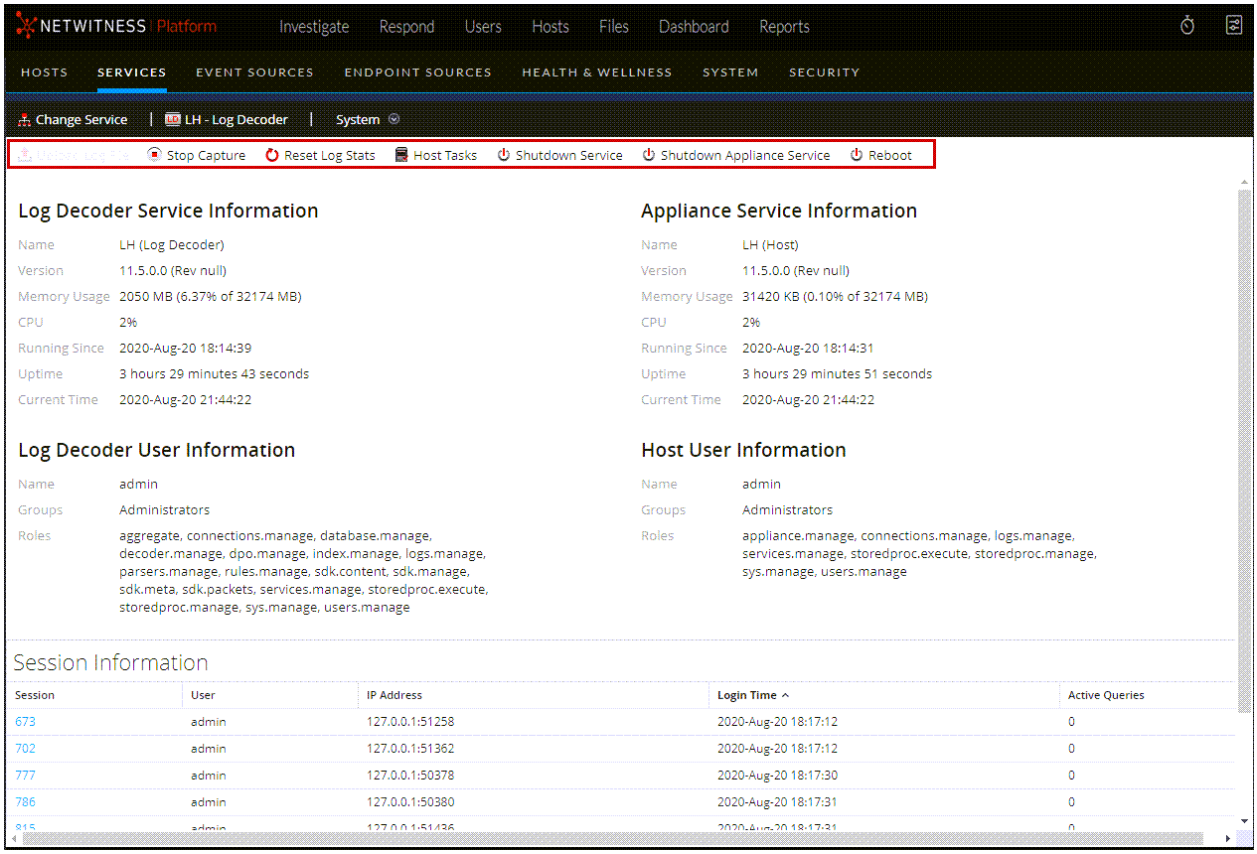
- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder
- "Services System View" in the *Hosts and Services Getting Started Guide*

Quick Look

This is an example of the Services System view for a Decoder.

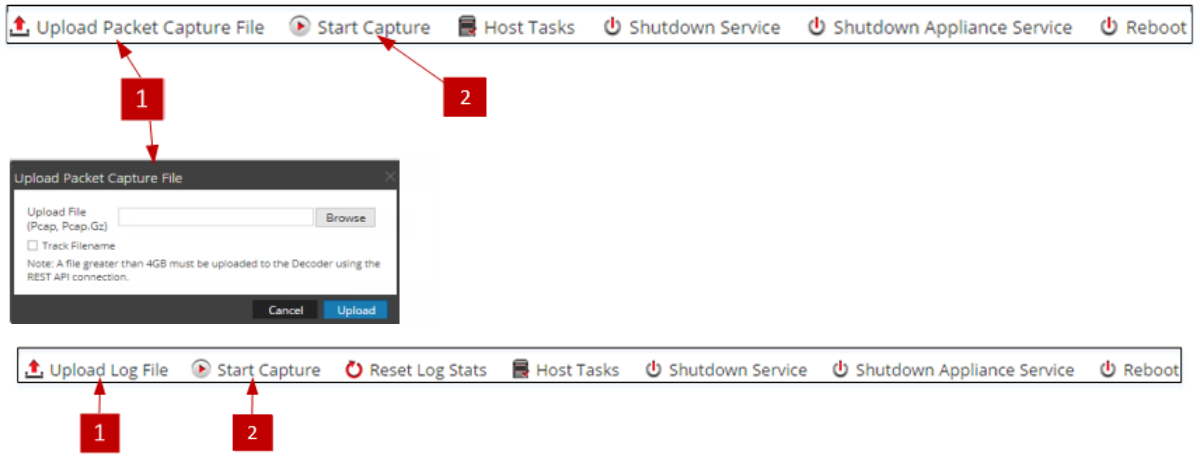


This is an example of the Services System view for a Log Decoder.



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.