

# ARISTA

## User Guide

## Arista Analytics

Version 8.6



<b>Headquarters</b>	<b>Support</b>	<b>Sales</b>
5453 Great America Parkway Santa Clara, CA 95054 USA +1-408-547-5500	+1-408-547-5502 +1-866-476-0000	+1-408-547-5501 +1-866-497-0000
<a href="http://www.arista.com/en/">www.arista.com/en/</a>	<a href="mailto:support@arista.com">support@arista.com</a>	<a href="mailto:sales@arista.com">sales@arista.com</a>

© Copyright 2024 Arista Networks, Inc. All rights reserved. The information contained herein is subject to change without notice. The trademarks, logos, and service marks ("Marks") displayed in this documentation are the property of Arista Networks in the United States and other countries. Use of the Marks is subject to the Arista Networks Terms of Use Policy, available at [www.arista.com/en/terms-of-use](http://www.arista.com/en/terms-of-use). Use of marks belonging to other parties is for informational purposes only.

# Contents

- Chapter 1: Arista Analytics Basic Operations..... 1**
  - 1.1 Overview..... 1
  - 1.2 Flows Dashboard..... 1
  - 1.3 Arista Analytics Fabric View..... 2
    - 1.3.1 Using Two-ring (by Production Switch) Pie Charts..... 2
    - 1.3.2 Filtering Information on a Dashboard..... 3
    - 1.3.3 Selecting the Time Range..... 3
    - 1.3.4 Using the Search Field..... 5
    - 1.3.5 Search Performance Limitations..... 5
  - 1.4 Using Discover Mode..... 5
  - 1.5 Managing Dashboards..... 6
  - 1.6 Custom Dashboards..... 7
  - 1.7 Mapping IP Address Blocks..... 11
  - 1.8 Mapping DHCP to OS..... 13
  - 1.9 Mapping Ports and Protocols..... 14
  - 1.10 SNMP Collector..... 15
  - 1.11 Mapping OUI to Hardware..... 16
  - 1.12 Topic Indexer on Arista Analytics..... 16
  
- Chapter 2: Production Network Monitoring..... 24**
  - 2.1 sFlow®..... 24
    - 2.1.1 sFlow and VXLAN..... 24
  - 2.2 NetFlow and IPFIX..... 25
    - 2.2.1 Consolidating Netflow V9/IPFIX records..... 27
    - 2.2.2 NetFlow and IPFIX Flow with Application Information..... 29
    - 2.2.3 NetFlow and sFlow Traffic Volume Upsampling..... 31
  - 2.3 TCPFlow..... 34
  - 2.4 Flows..... 35
  - 2.5 Filters & Flows..... 36
  - 2.6 ARP..... 37
  - 2.7 DHCP..... 38
  - 2.8 DNS..... 39
  - 2.9 ICMP..... 40
  
- Chapter 3: Using the DMF Recorder Node with Analytics..... 42**
  - 3.1 Overview..... 42
  - 3.2 General Operation..... 42
  - 3.3 Using Recorder with Analytics..... 44
  - 3.4 Analyzing SIP and RTP for DMF Analytics..... 45
  
- Chapter 4: Managing the NetFlow Dashboard..... 47**
  - 4.1 NetFlow Optimization..... 47
  - 4.2 Viewing Filter Interface Information on the NetFlow Dashboard..... 48
    - 4.2.1 Displaying Filter Interface Names..... 48
    - 4.2.2 NetFlow Traffic Coming from Third-party Devices..... 50
  - 4.3 Displaying Flows with Out-Discards..... 53

---

<b>Chapter 5: Advanced Feature Dashboard.....</b>	<b>54</b>
5.1 Latency Differ and Drop Differ Dashboard.....	54
<b>Chapter 6: Monitoring DMF Network Health.....</b>	<b>66</b>
6.1 DMF Network Tab.....	66
6.2 Policy Statistics Dashboard.....	66
6.3 Interface Statistics.....	68
6.4 SN (Service Node) Statistics.....	68
6.5 Events.....	69
<b>Chapter 7: Monitoring Users and Software Running on the Network.....</b>	<b>71</b>
7.1 IP Addresses.....	71
7.1.1 Source and Destination Addresses.....	71
7.1.2 Unauthorized IP Destinations.....	72
7.2 Geographic Location.....	73
7.3 Software Running in the Network.....	74
7.3.1 Top Talkers Using Well-known Layer-4 Ports.....	74
7.3.2 Associating Applications with User-defined Layer4 Ports.....	76
7.3.3 Software Running on Hosts.....	77
7.3.4 Tools Receiving Traffic.....	78
7.4 User Activity.....	81
7.4.1 User Sessions.....	81
7.4.2 New Network Users.....	83
7.4.3 Unauthorized Intranet Activity.....	84
7.5 Monitoring Active Directory Users.....	85
<b>Chapter 8: Monitoring Network Performance and Events.....</b>	<b>86</b>
8.1 Interfaces Sending or Receiving Traffic.....	86
8.2 Anomalies.....	88
8.3 Application Data Management.....	89
8.4 WAN Link Optimization.....	90
8.5 Machine Learning.....	92
<b>Chapter 9: Backup and Restore.....</b>	<b>96</b>
9.1 Elasticsearch Snapshot and Restore.....	96
9.2 Import and Export of Saved Objects.....	97
9.2.1 Exporting Saved Objects.....	97
9.2.2 Importing Saved Objects.....	98
9.3 Import and Export of Watchers.....	99
9.3.1 Exporting Watchers.....	100
9.3.2 Importing Watchers.....	100
9.4 Import and Export of Machine Learning Jobs.....	102
9.4.1 Exporting Machine Learning Jobs.....	102
9.4.2 Importing Machine Learning Jobs.....	103
<b>Chapter 10: Using TACACS+ and RADIUS to Control Access to the Arista Analytics CLI.....</b>	<b>105</b>
10.1 Using AAA Services with Arista Analytics.....	105

- 10.1.1 DMF TACACS+ Configuration..... 106
- 10.2 Adding a TACACS+ Server..... 107
- 10.3 Setting up a TACACS+ Server..... 108
  - 10.3.1 Using the Same Credentials for the Analytics Node and Other Devices..... 108
  - 10.3.2 RBAC-based Configuration for Non-default Group User..... 109
- 10.4 Using RADIUS for Managing Access to the Arista Analytics Node..... 109
  - 10.4.1 Adding a RADIUS Server..... 110
  - 10.4.2 Setting up a FreeRADIUS Server..... 110
  
- Appendix A: Creating Watcher Alerts for Machine Learning jobs..... 112**
  - A.1 Watcher Alert Workaround..... 112
  - A.2 Email Alerts and Remote Syslog Server..... 119
  - A.3 Enabling Secure Email Alerts through SMTP Setting..... 125
  
- Appendix B: References..... 127**
  - B.1 Related Documents..... 127



## Arista Analytics Basic Operations

---

This chapter uses Arista Analytics to monitor and analyze traffic and events in the monitoring fabric and the DANZ Monitoring Fabric controller. This chapter includes the following sections:

- [Overview](#)
- [Flows Dashboard](#)
- [Arista Analytics Fabric View](#)
- [Using Discover Mode](#)
- [Managing Dashboards](#)
- [Custom Dashboards](#)
- [Mapping IP Address Blocks](#)
- [Mapping DHCP to OS](#)
- [Mapping Ports and Protocols](#)
- [SNMP Collector](#)
- [Mapping OUI to Hardware](#)
- [Topic Indexer on Arista Analytics](#)

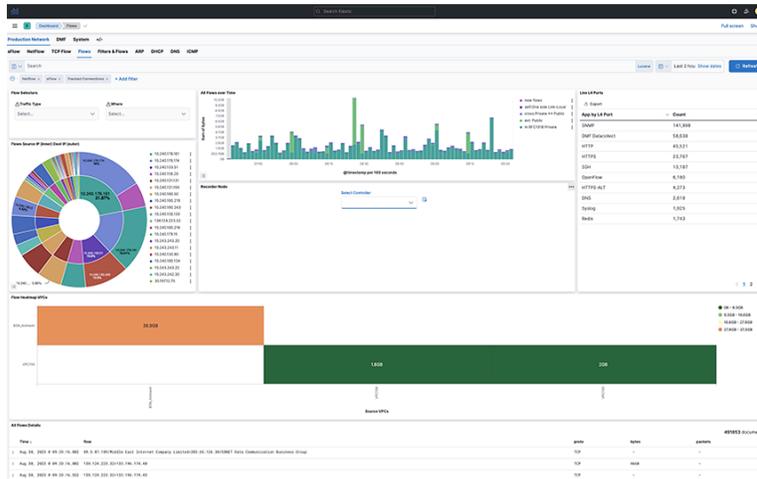
### 1.1 Overview

Arista Analytics provides a non-proprietary extensible UI that integrates DMF Recorder Nodes, DMF Service Nodes, and the DANZ Monitoring Fabric controlled using an SDN Controller. The system has an extensive library of visualization components and analytics to compose new dashboards and answer further questions as they arise. The Arista Analytics node/cluster answers questions that would otherwise require specialized applications, such as Application Data Management (ADM) or Intrusion Protection Management (IPM). The Analytics node/cluster creates a document for each packet received. It adds metadata regarding the context, including the time and the receiving interface, which Elasticsearch can use to search the resulting documents quickly and efficiently.

## 1.2 Flows Dashboard

The following figure shows the Flows dashboard when accessing Arista Analytics.

**Figure 1-1: Production Network > Flows Dashboard**



The left panel provides the following options to access Arista Analytics features:

- **Fabric:** The home page for Analytics provides a series of tabs and sub-tabs.
- **Controller:** Opens the DANZ Monitoring Fabric GUI on the Controller identified using the **System > DMF Controller** option.
- **Discover:** Use predefined indices to filter and display specific events.
- **Visualize:** Customize the graphics displays provided by Arista Analytics.
- **Dashboard:** Displays dashboards for DANZ Monitoring Fabric events.
- **Timelion:** Displays events and other results according to time series.

The Kibana documentation documents the Analytics GUI, and most of its features and operations based on Elasticsearch are available at the following URL:

<https://www.elastic.co/guide/en/kibana/7.2/index.html>

**Kibana 7.2** is the version used for **Arista Analytics version 7.3**.

## 1.3 Arista Analytics Fabric View

The Arista Analytics Fabric view displays in the following three tabs:

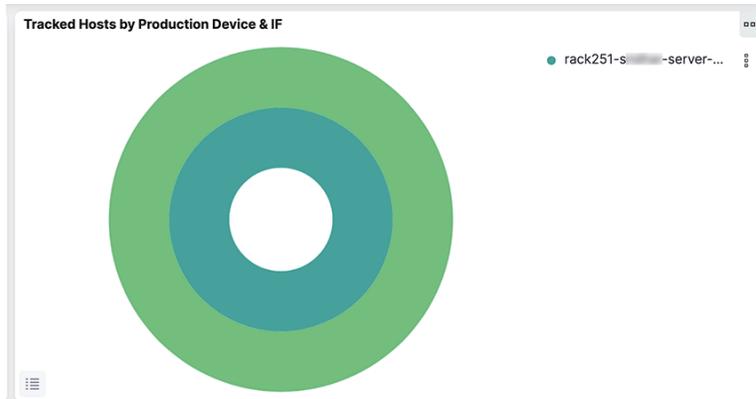
- **Production Network:** View information about the production network.
- **DMF Network:** View information about the monitoring network.
- **System:** Manage system configuration settings.

Each page contains panels with different functions and features. The network panels provide visualizations, such as pie charts, line graphs, or other graphic displays that reflect the current dashboard contents based on the specific query. The bottom panel lists all the events that match the current query. A pop-up window provides additional details about the selection when mousing over a panel.

### 1.3.1 Using Two-ring (by Production Switch) Pie Charts

Pie charts that display information by the production switch have an inner and outer ring, as shown in the following example.

**Figure 1-2: Two-ring Pie Chart**



When a second ring appears in a pie chart, click any segment in the inner ring, and the outer ring provides a summary of information about the selected segment.

For example, in the **Tracked Hosts by Production Device & IF** pie chart, the outer ring shows hosts tracked on each interface, while the inner ring summarizes the tracked hosts on each switch. Clicking on a segment for a specific switch on the inner ring filters the outer ring to show the tracked hosts for the interfaces on the selected switch.

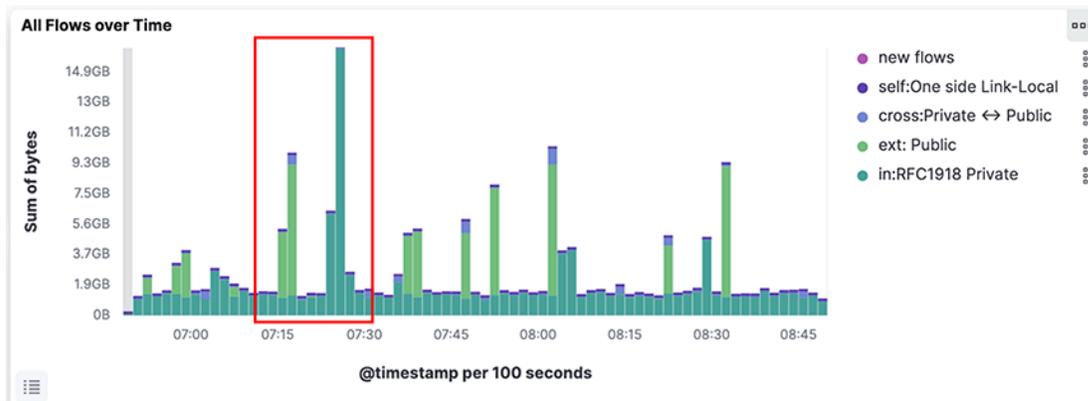
### 1.3.2 Filtering Information on a Dashboard

Filter the events displayed on a dashboard by choosing an area on the dashboard. This action limits the information displayed on the dashboard to events similar to those selected. Click any pie chart slice to limit the display to the specific activity chosen. To change the color assigned to a specific protocol or other object, click the label on the list to the right of the chart.

### 1.3.3 Selecting the Time Range

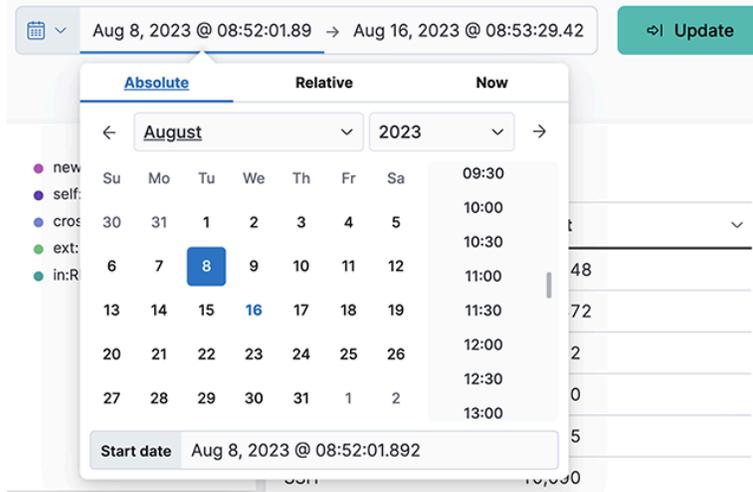
To restrict the current content to events occurring in a specific period, click the mouse and drag it to surround the area on a time visualization, such as the Flows Over Time.

**Figure 1-3: Selecting the Time Range**



To select the time range or to change the default refresh rate, click the **Time Range** control in the upper right corner. The system displays the following dashboard.

**Figure 1-4: Time Range Control**

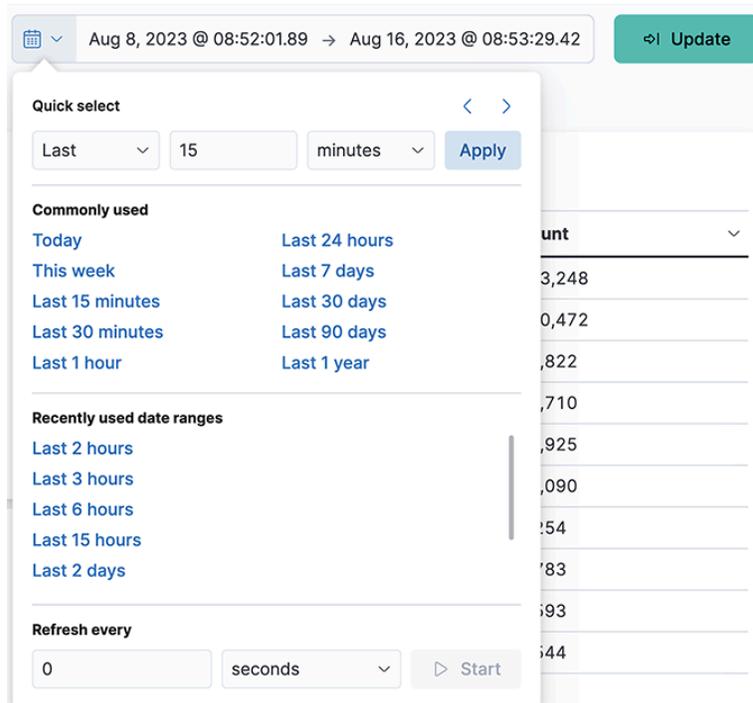


This dialog provides the following options for setting the time range:

- **Quick:** Simple settings, such as Today, Last 1 hour, etc.
- **Relative:** Time offsets from a specific time, including the current time.
- **Absolute:** Set a range based on date and time.
- **Recent:** Provides a list of recently used ranges that you can reuse.

Select the range from the options provided, and the panels and displays update to reflect the new date and time range. To change the auto-refresh rate, click the **Auto-refresh** control. The system displays the following dashboard.

**Figure 1-5: Change Auto Refresh Rate**

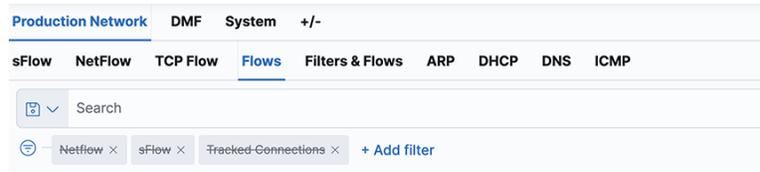


Select the refresh interval from the options provided. Click **Start** to disable the auto-refresh function.

### 1.3.4 Using the Search Field

The search field at the top of the dashboard filters the current displays by any text or numbers you type into the field.

**Figure 1-6: Search Field**



The green bars under the **Search** field show the currently applied filters. When the pointer is over a green bar, it displays icons that let you control the filter.

- **Enable/Disable filter**
- **Pin/Unpin filter**
- **Exclude/Include matches**
- **Remove filter**
- **Edit filter**

The **Action** option in the upper right corner applies these actions to all the currently applied filters.

Click a segment on a pie chart for the appropriate filter; it automatically inserts into the **Search** field. To undo the filter, click the **Remove** filter icon.

To filter the information in the displays, enter the characters to filter the display in the search field. For example, for entering the first part of an IP address, it updates the displays to show only those IP addresses that match the characters entered. The following are some of the most helpful search filters:

- IP address
- Host name (requires DNS services)
- Protocol, for example, HTTP, HTTPS, ICMP, and so forth
- DMF interface name

To define complex queries using field names, which can be seen by scrolling and clicking on an event row.

For example, on the sFlow<sup>®</sup> dashboard, the query `proto : TCP AND tags : ext` displays all externally bound TCP traffic. OR NOT ( ) are also permitted in the expression. For more details about the supported search syntax, refer to the following URL: <https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-query-string-query.html#query-string-syntax>.

### 1.3.5 Search Performance Limitations

Refrain from executing a general query for 7 or 30 days. You should select a 7 or 30-day query with specific criteria, like querying a specific flow, filter interface, or DNS server.

To query NetFlow or sFlow for more extended periods, use the **FLOW** dashboard to determine the trend and then do a specific query, such as querying a specific flow or time, on the Netflow or sFlow dashboard.

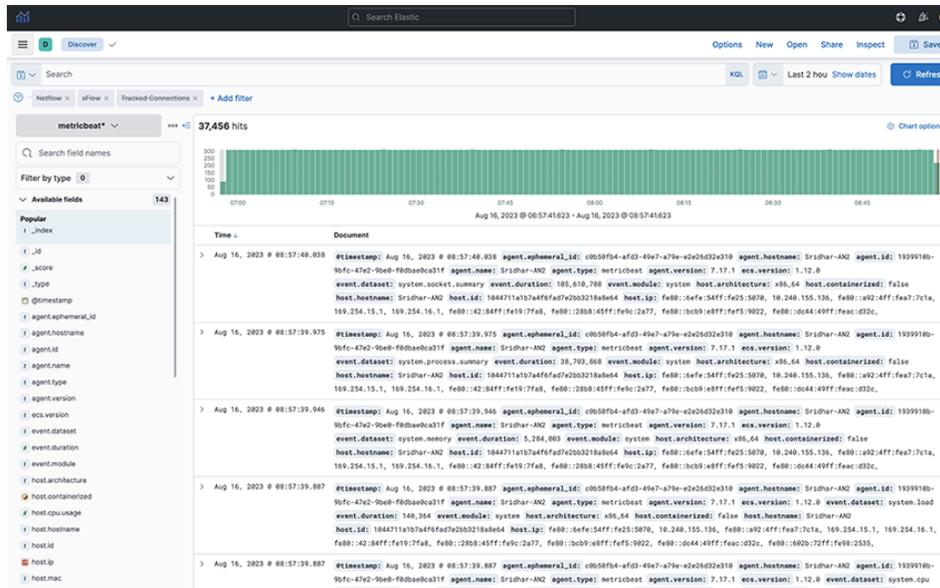
For a great deal of NetFlow traffic, select one Analytics node only for NetFlow and another for other Analytics traffic.

\* sFlow<sup>®</sup> is a registered trademark of Inmon Corp.

## 1.4 Using Discover Mode

Select the **Discover** option in the left panel of the **Analytics** window, the system will display the following page.

Figure 1-7: Discover Mode

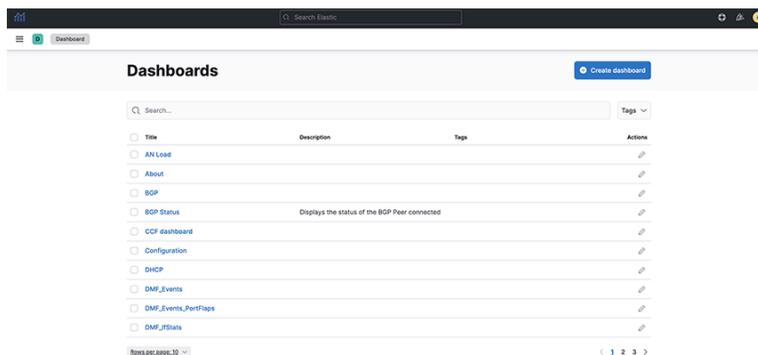


Use **Discover** mode to see the indices in the ElasticSearch database and identify the available data.

## 1.5 Managing Dashboards

Select the **Dashboards** option from the left panel on the **Analytics** window to manage dashboards. The system displays the following page.

Figure 1-8: Dashboard Mode



Refer to the Kibana documentation for details about creating and managing dashboards. <https://www.elastic.co/guide/en/kibana/7.13/index.html>



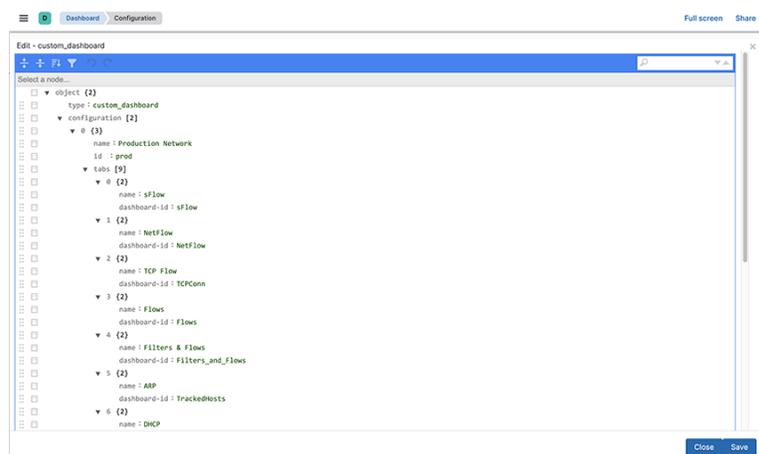
**Note: Recommended Best Practices** - Use the naming convention that suits your environment while creating a dashboard or saved objects. For example, select a prefix to identify the dashboard content, and then use the body of the dashboard name to determine the type of dashboard. For instance, in the above screenshots, it uses a naming pattern, prefixed with "ARISTA" and specifying type: dashboard allows a manageable set of things to appear to click or select all individually. Furthermore,

exporting individual dashboards based on their type is a more appropriate option for easy tracking as modifications to a dashboard. Your dashboards should use only visualizations and searches you create for upgrades; do not depend on default objects that might change in the upgrade.

## 1.6 Custom Dashboards

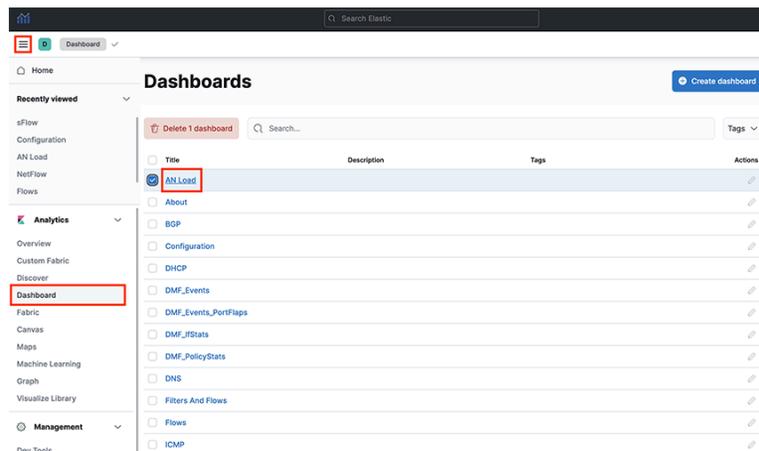
Add or insert the custom dashboard by selecting the **Dashboards** option from the left panel on the **Analytics** window. The system displays the following page, which is the default dashboard:

**Figure 1-9: Default Dashboard Mode**



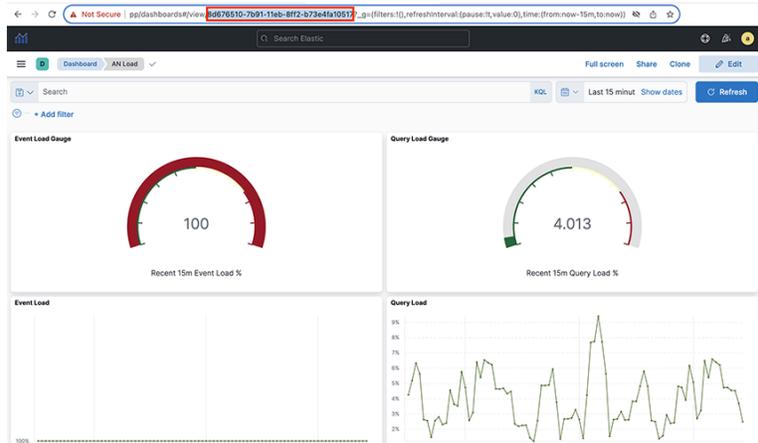
In the default dashboard, select the option to customize per your requirements.

**Figure 1-10: Search for Dashboard**



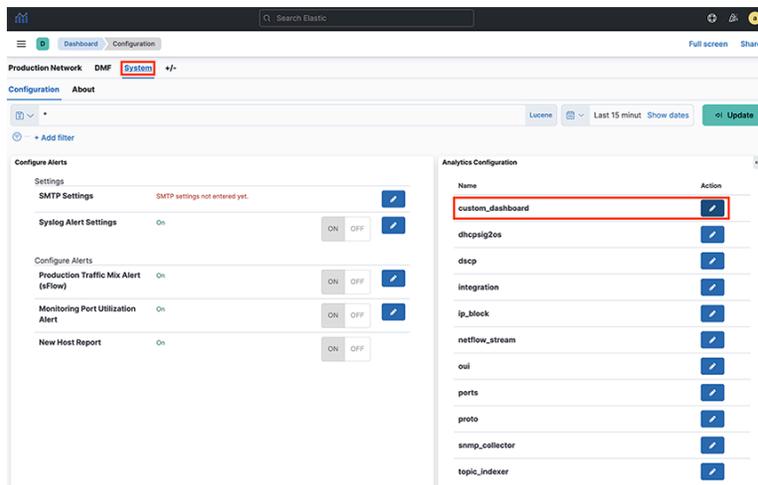
For the customization of the option on the dashboard, copy its ID as following.

**Figure 1-11: Select the option and copy the ID**

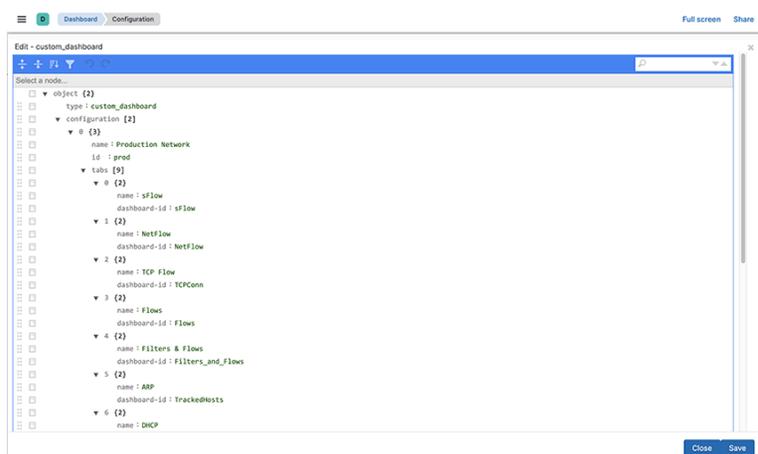


**Note:** Insert the ID into the dashboard in the same way as captured from the bar to work.

**Figure 1-12: Setting custom Dashboard**



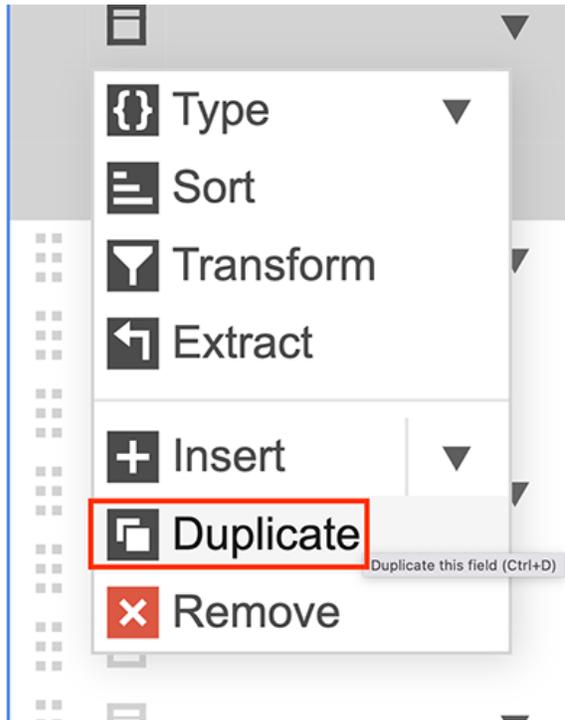
**Figure 1-13: Default Dashboard configuration**



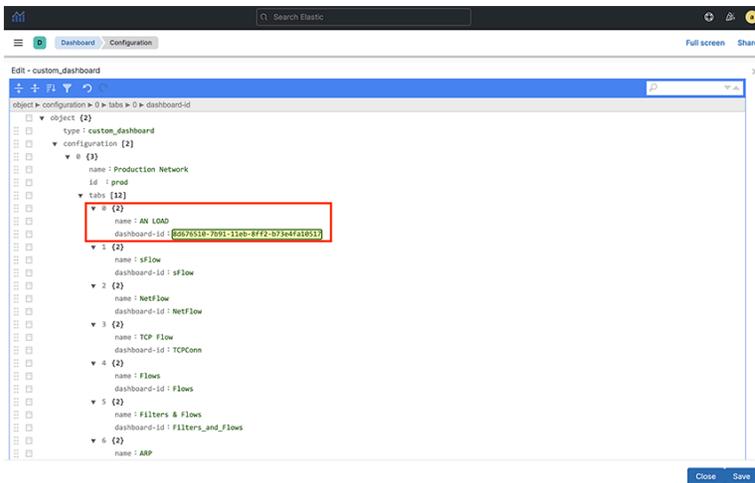


Select the **Duplicate** tab for the duplicate entries.

**Figure 1-15: Duplicate the tab**

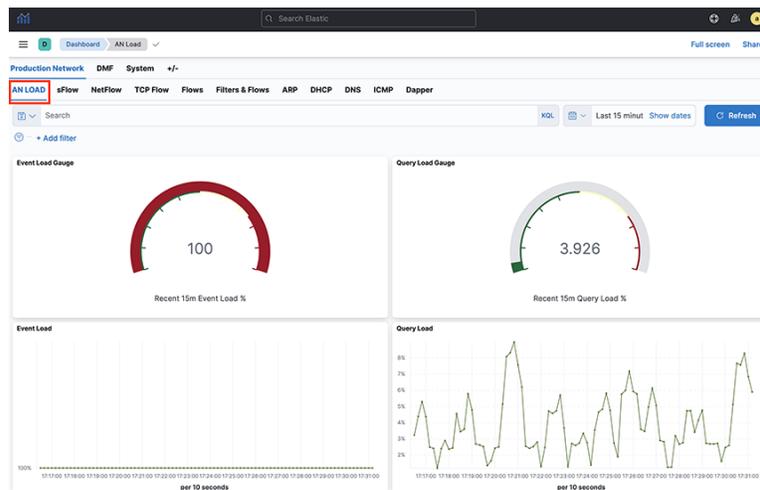


**Figure 1-16: Insert the name tag ID**



Now, the dashboard shows the customization of the option selected by the user.

**Figure 1-17: Selected option for the user**



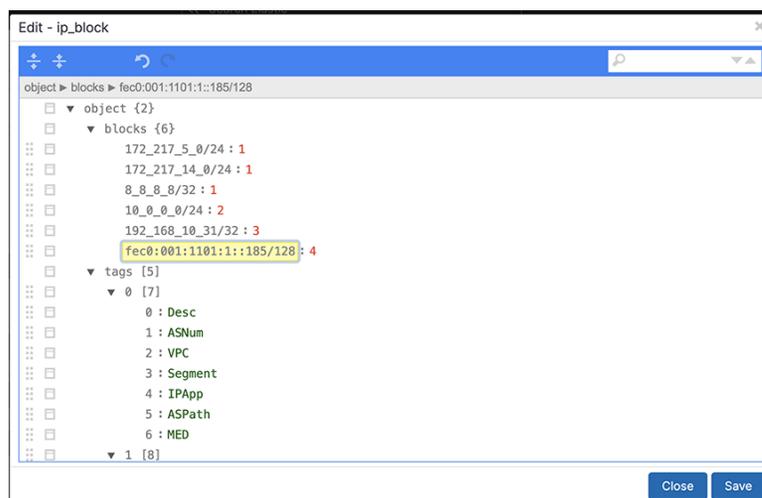
## 1.7 Mapping IP Address Blocks

Map an IP address or a range of addresses to a description, which searches for description text instead of the IP address. This feature identifies a specific group or organization sending or receiving traffic.

Complete the following steps to assign a single IP address or a block of IP addresses to a tool, group, or organization.

1. Select **System > Configuration** and click the **Edit** control to the left of the IP Block section.

**Figure 1-18: Edit IP Blocks**

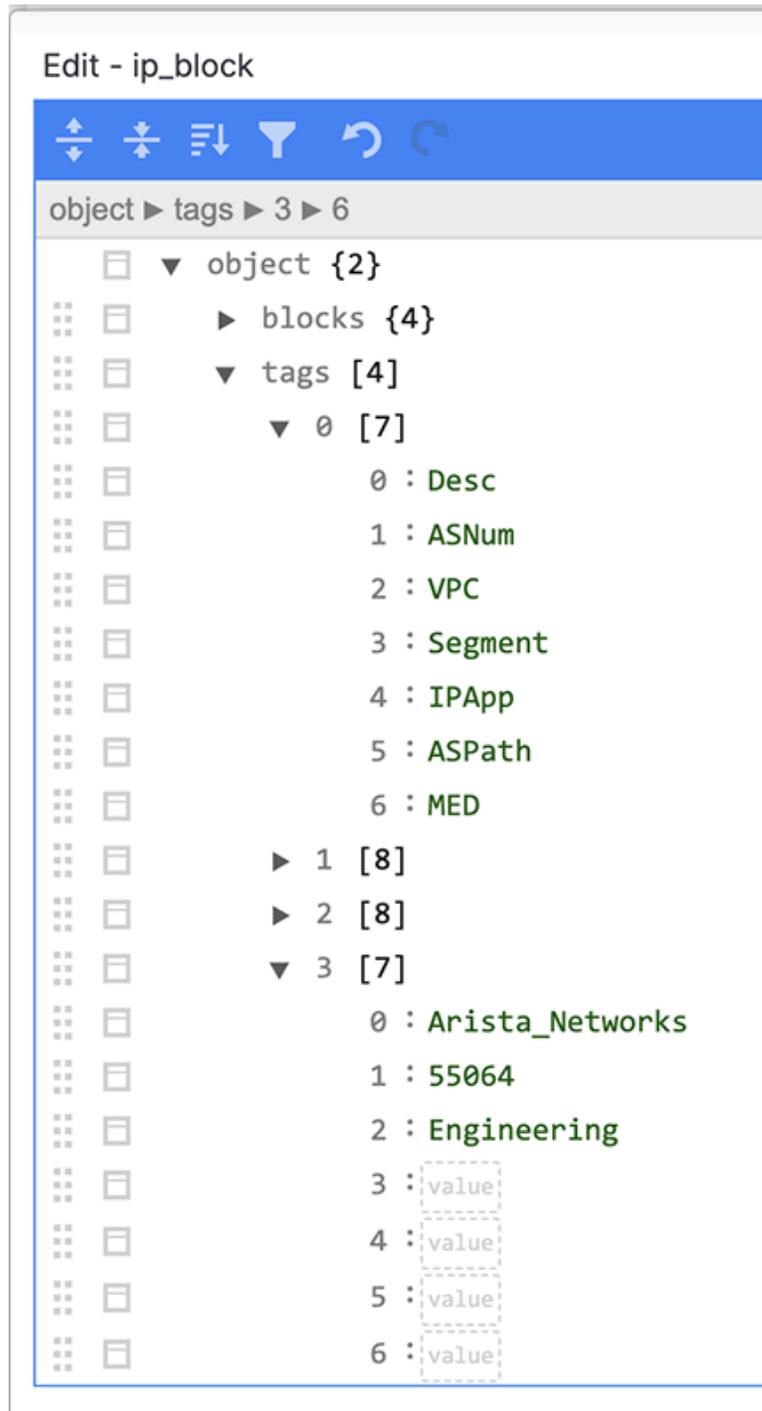


2. Copy an existing block by clicking on any square box along the left and select **Duplicate** from the pop-up menu.

The duplicated block will be appended to the existing block list and assigned the next numerical sequence identifier.

3. Scroll down to the end of the tags section to the numerical identifier assigned to the new block.

**Figure 1-19: Key Value Pairs**



It automatically copies the first four keys.. The purpose of each of these default keys is as follows.

- **Desc:** A short descriptive text entry.
- **ASNum:** Automatically populated with the BGP Autonomous Systems (AS) numbers for well-known networks.
- **VPC:** Virtual Private Cloud (tenant), automatically populated with the VPCs used in an integrated Converged Cloud Fabric network.
- **Segment:** Network segment within a Converged Cloud Fabric VPC.

To identify a user, application, tool, group, or organization, use the **Desc** key. You can leave the other fields blank.

4. Type a value for the Desc key in double quotation marks ("").
5. (Optional) To define an additional key, select any key and choose **Duplicate** from the pop-up menu. Then, type over the existing value with the correct value for the new key.

Existing dashboards use the default keys. The customized dashboards can use added key pairs. The fifth and sixth keys can be custom.

These keys are added to the flow for the source and destination IPv4 address. For example, the source description would be **sDesc** and the destination description would be **dDesc**.



**Note:** Remember to match values in the same order as the corresponding key positions.

## 1.8 Mapping DHCP to OS

DHCP signatures can map to known operating systems. These unique signatures are from **#ngerbank.org**. As shown in the following image, several two-digit numbers are assumed signatures of each OS (derived from **#ngerbank.org**).

Figure 1-20: Unique OS Signatures from fingerbank.org

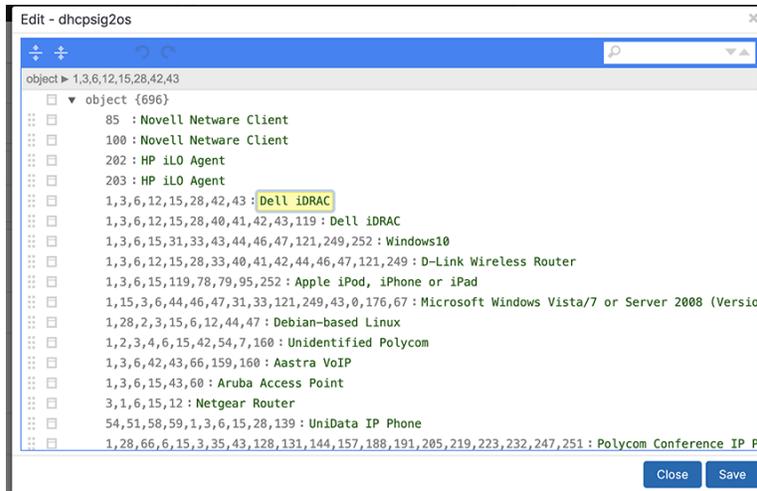
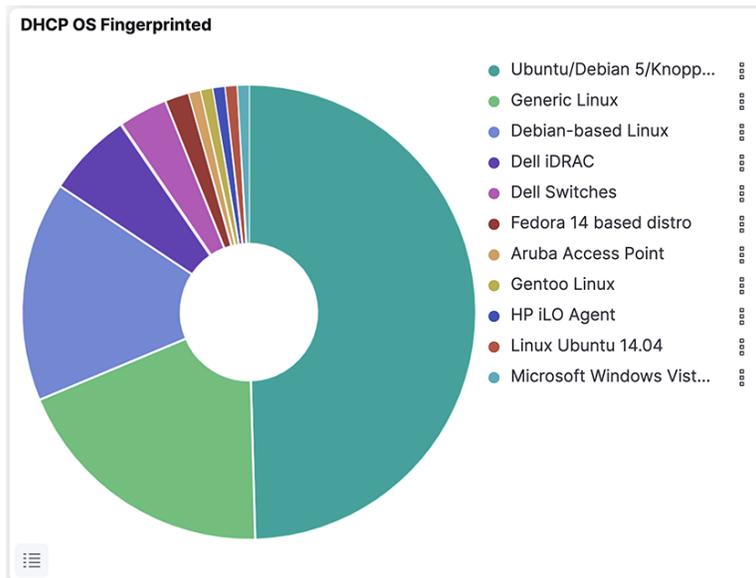


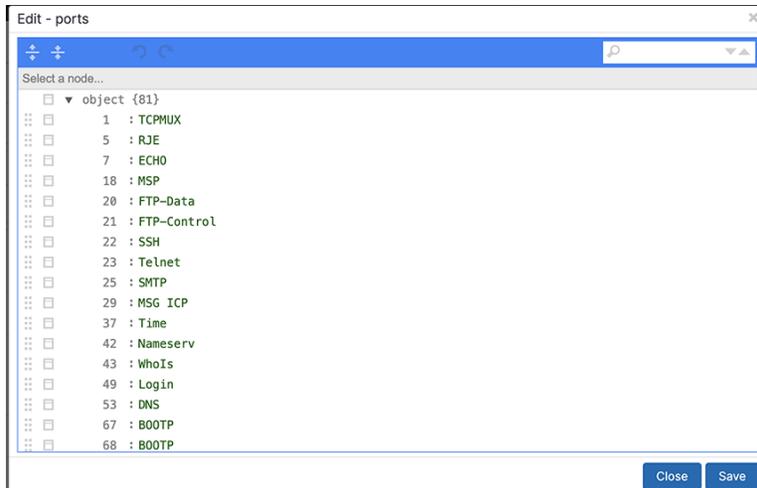
Figure 1-21: OS Information Received through DHCP Signatures



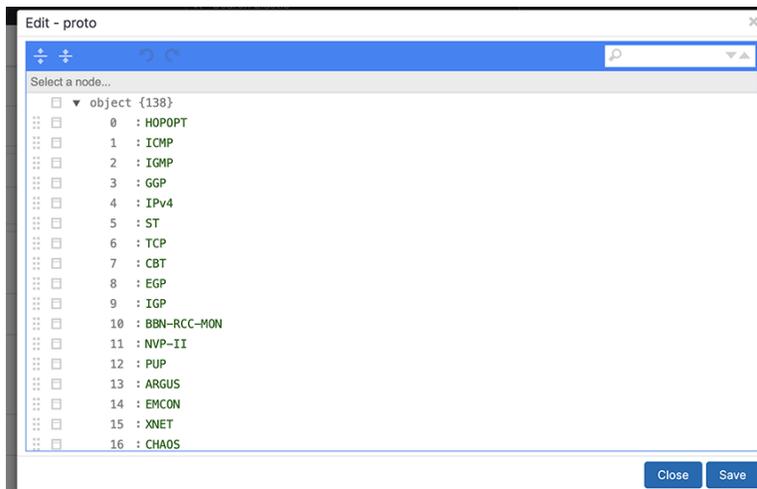
## 1.9 Mapping Ports and Protocols

The Analytics Node maps typically used ports for their L4 applications and protocols. These protocols and ports can also be user-defined for custom application ports and custom protocols.

**Figure 1-22: Edit Ports**



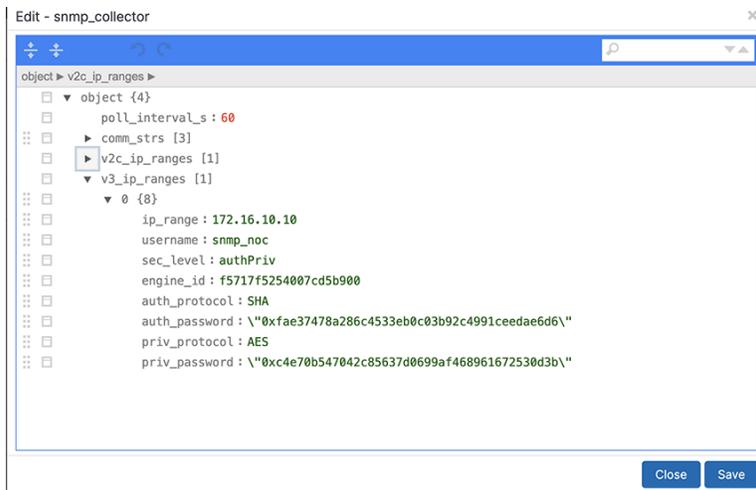
**Figure 1-23: Edit Protocols**



## 1.10 SNMP Collector

SNMP collectors facilitate third-party NetFlow/IPFIX sources. The Analytics Node supports both SNMPv2 and SNMPv3.

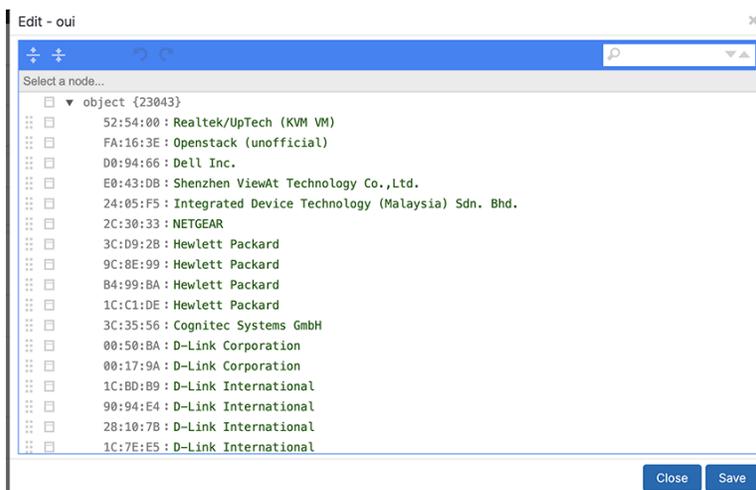
Figure 1-24: SNMP Collector



## 1.11 Mapping OUI to Hardware

Map ARP Organizational Unique Identifiers (OUIs) for various hardware vendors.

Figure 1-25: OUIs of Various Hardware Vendors



## 1.12 Topic Indexer on Arista Analytics

### Description

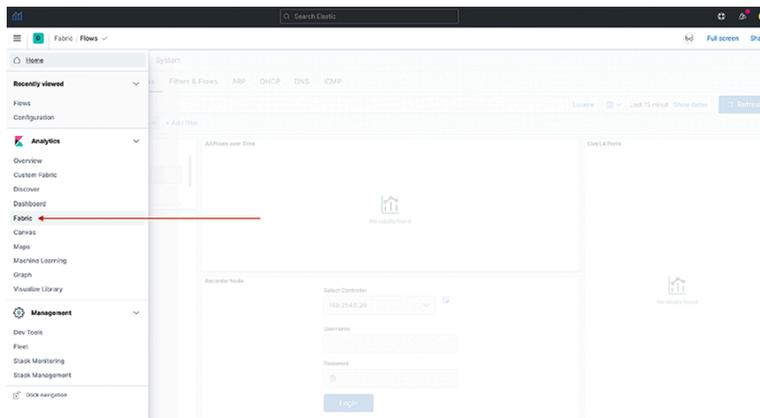
The Analytics Node (AN) incorporates a feature known as topic\_indexer, designed to facilitate data ingestion from customer Kafka topics and its subsequent storage into Elasticsearch indices.

This process involves modifying field names and specifying the supported timestamp field during the ingestion phase. The renaming of field names enables the creation of dashboards used to visualize data across multiple streams, including DNS and Netflow.

The resulting indices can then be leveraged as searchable indices within the Kibana user interface, providing customers with enhanced search capabilities.

### Implementation Details

- Configure a stream job using `topic_indexer`. Access the setting via the Kibana dashboard in the analytics node.
- Locate the `topic_indexer` configuration on the Fabric Dashboard: **Analytics > Fabric > System > Analytics Configuration**, as shown in the following screenshots.
- **Figure 1-26: Analytics > Fabric**

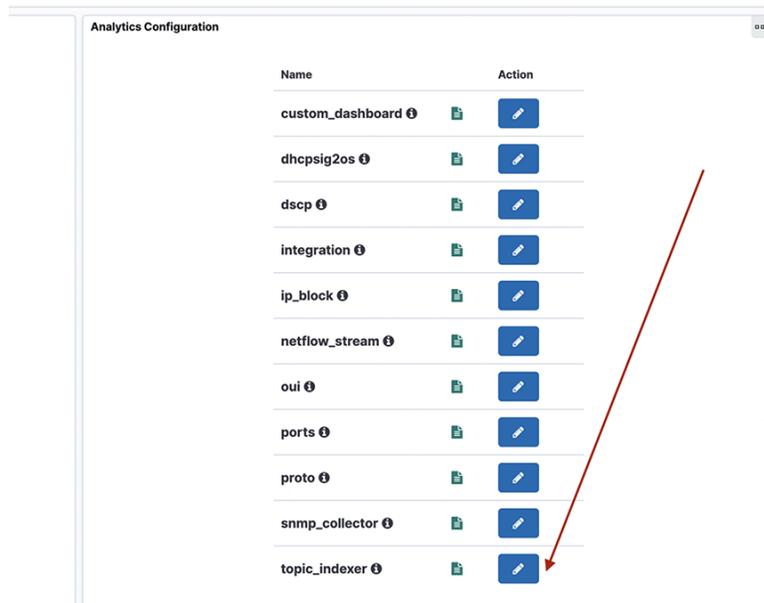


- Another view:

### Figure 1-27: System > Analytics Configuration



- The design section shows the configuration for a topic
- **Figure 1-28: Node selection**



The screenshot shows the 'Analytics Configuration' page in Kibana. It features a table with two columns: 'Name' and 'Action'. Each row represents a configuration node, with a green document icon and a blue edit icon. A red arrow points to the 'topic\_indexer' node at the bottom of the list.

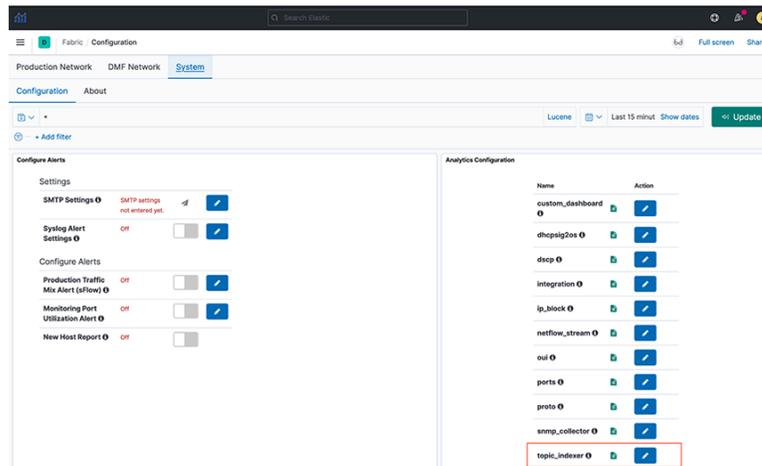
Name	Action
custom_dashboard ⓘ	
dhcpsig2os ⓘ	
dscp ⓘ	
integration ⓘ	
ip_block ⓘ	
netflow_stream ⓘ	
oui ⓘ	
ports ⓘ	
proto ⓘ	
snmp_collector ⓘ	
topic_indexer ⓘ	

**Configuration**

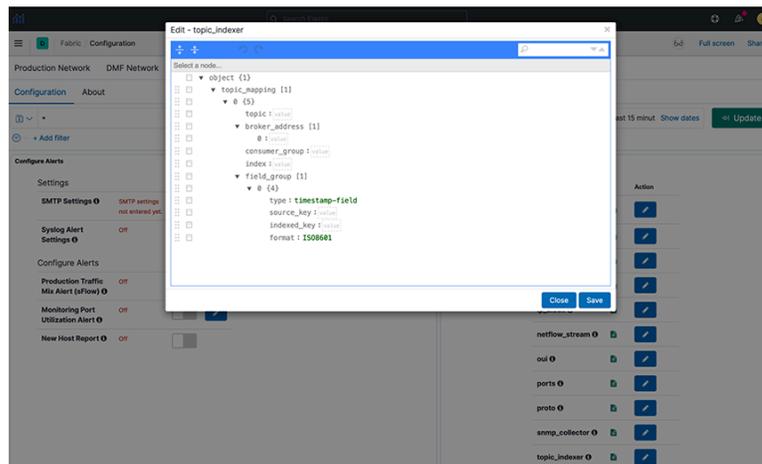
**Kibana Configuration**

- To perform the **topic\_indexer** configuration, select the **System > Configuration > Fabric** page and open the **Analytics Configuration** panel:

**Figure 1-29: System > Configuration**



- Figure 1-30: Topic\_indexer configuration**



## Field Details

Each topic maps in JSON with the following fields:

- topic:** Kafka topic name; type string and is a mandatory field.
- broker\_address:** Broker address(es), this is of type array; this will be of format **[IPv4]hostname:Port number** and is a mandatory field.
- consumer\_group:** This is an optional field; however, there is always a consumer group if not specified explicitly in the configuration. It is **topic\_name + index\_name**. Setting this field is particularly useful when ingesting multi-partitioned topics from the client's end.
- index:** A dedicated index name for the topic; type string. In Elastic Search (ES), it is created as **topic\_indexer\_<index\_name>** and is a mandatory field.
- field\_group:** An optional JSON field mapping to specify any column rename/format transformations. It specifies format for modifications to incoming data.
- type:** To set timestamp field as the type.
- source\_key:** The source field name in the incoming data.
- indexed\_key:** The name of the destination field inserted in the outgoing ES index.

---

The **indexed\_key** may be a `@timestamp` field of an ES index. If you do not specify a `@timestamp` field, **topic\_indexer** automatically picks the time the message was received as the `@timestamp` of that message.

- **format:** Data format for the field (ISO8601).

## Standards and Requirements

Input fields naming convention:

- Kafka allows all ASCII Alphanumeric characters, periods, underscores, and hyphens to name the topic. In **topic\_indexer**, legal characters include: **a-z0-9\\_\|\-**
- Note that the only restriction **topic\_indexer** has is on capitalizing topic names. **topic\_indexer** does not support case-sensitive names. By default, **topic\_indexer** treats the name as a lowercase topic. Hence, topic names should be lowercase only.
- All numeric names are also invalid field text.



**Note:** These conventions are valid for all other input types as well.

## Examples of names:

### Valid text:

- my-topic-name
- my\_topic\_name
- itlabs.mytopic.name
- topic123
- 123topic
- my-index-name

### Invalid text:

- myTopicName
- ITLabs-Website-Tracker
- 12435
- MY-Index-name

## Broker Address Format:

- A broker address in Kafka comprises two values: IPv4 address and Port Number.
- When entering the broker address, use the format: **IPv4:PORT**.

## Application Scenario

### Querying Across DataStream using runtime-fields

Use runtime fields when making complex changes beyond simply renaming a field, such as converting it from a string type to an IP address. After every change to a runtime field, issue a

```
POST <stream-name>/_rollover
```



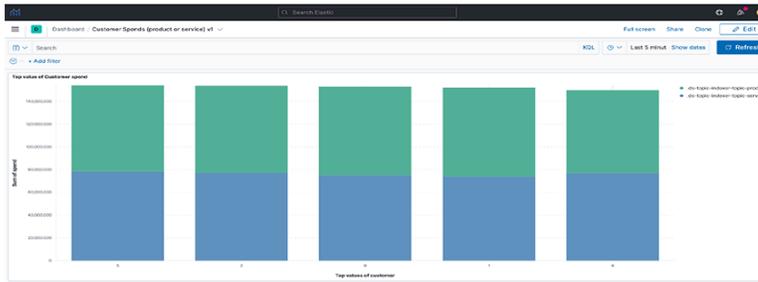
**Note:** These changes are not persistent. Reapply is a must after any restart of AN.

### Use Case:



- **Step 4.** Finally, create a visualization using this common field, for example, **Customer**. The illustration below shows the Top 5 customers with the highest spending across products and services.

**Figure 1-34: Visualization**



## Syslog Messages

The **topic\_indexer** logs are stored in `/var/log/analytics/` folder and are accessed using the following commands.

```
an> debug bash
admin@an$ cd /var/log/analytics/
admin@an:/var/log/analytics$
admin@an:/var/log/analytics$ ls -ls topic_indexer.log
67832 -rw-rwxr-- 1 remoteuser root 69453632 Apr 27 11:05 topic_indexer.log
```

## Troubleshooting

Below are some of the commonly known issues and their troubleshooting scenarios:

### 1. The Save button in the topic\_indexer config is disabled.

When editing the configurations of **topic\_indexer** in the Kibana User interface, default validations appear to ensure the correctness of the values entered in the fields. Specific standards and requirements are associated when filling in the config for **topic\_indexer** as stated in the above section linked: [Topic Indexer on Arista Analytics](#) . As illustrated below, it may encounter validation errors when entering an invalid value in the configuration field. [Topic Indexer on Arista Analytics](#)

**Figure 1-35: Validation errors**



In such an event, the edited configuration will not save. Therefore, before saving the configuration, validate the fields and ensure there is no visible validation error in the **topic\_indexer** configuration editor.

### 2. The index for the topic\_indexer is not created.

After entering the correct fields in the **topic\_indexer** configuration, the **topic\_indexer** service will start to read the Kafka topic as documented in the configuration and load its data into the ElasticSearch index entered by the index field. The name of the index is prefixed by **topic\_indexer\_**.

There is a wait time of several minutes before the index is created and loaded with the data from the Kafka topic. In the event the index is not created, or there is no index shown with the name **topic\_indexer\_<index\_name>** value, Arista Networks recommends using the following troubleshooting steps:

- a. Check the configurations entered in the **topic\_indexer** editor once again to see whether the spellings for the topic name, broker address configuration, and index name are correct.
- b. Verify the broker address and the port for the Kafka topic are open on the firewall. Kafka has a concept of listeners and **advertised.listeners**. Validate if the **advertised.listeners** are entered correctly into the configuration. Review the following links for more details:
  1. [Kafka 3.5 Documentation](#)
  2. [Kafka Listeners – Explained | Confluent](#)
- c. If all the above steps are correct, check now for the logs in the Analytics Node for the **topic\_indexer**.

**Steps to reach the topic\_indexer.log file in AN node:**

1. Secure remote access into the AN using the command line: `ssh <user>@<an-ip>`
  2. Enter the password for the designated user.
  3. Enter the command `debug bash` to enter into debug mode.
  4. Use the sudo user role when entering the AN node: hence the `sudo su` command.
  5. **topic\_indexer** logs reside in the following path: `/var/log/analytics/topic_indexer.log`
  6. Since this log file can be more extensive, you should use the tail command.
  7. Validate if the log file shows any visible errors related to the index not being created.
  8. Report any unknown issues.
3. **Data is not indexed as per the configuration.**
  4. **Data ingestion is paused.**

When experiencing issues 3 or 4 (described above), use the **topic\_indexer** log file to validate the problem.

5. **The index pattern for the topic\_indexer is missing.**

In the Kibana UI, it creates a default **topic\_indexer\_\*** index pattern. If this pattern or a pattern to fetch the dedicated index for a topic is missing, create it using the Kibana UI as described in the following link:

[Create an index pattern | Kibana Guide \[7.17\] | Elastic](#)

# Production Network Monitoring

---

This chapter describes the dashboards provided on the Production Network tab, which shows traffic and events on the production network interfaces connected to the DANZ Monitoring Fabric. This chapter includes the following sections:

- [sFlow®](#)
- [NetFlow and IPFIX](#)
- [TCPFlow](#)
- [Flows](#)
- [Filters & Flows](#)
- [ARP](#)
- [DHCP](#)
- [DNS](#)
- [ICMP](#)

## 2.1 sFlow®

Click the **Fabric** option; it displays the sFlow® dashboard by default. It summarizes information from the sFlow messages sent to the Arista Analytics server from the DANZ Monitoring Fabric controller or other sFlow agents. This dashboard provides the following panels:

- Top Sources
- Source Port
- Top Destinations
- Destination Port
- Traffic over time
- Flow by Filter Interface
- Flow by Device & IF
- Count sFlow vs. Last Wk
- Flow QoS PHB
- Flow Source
- Flow Destination
- sFlow MTU Distribution
- Flows by Time

### 2.1.1 sFlow and VXLAN

The sFlow dashboard shows both outer and inner flows of VXLAN packets based on the VNI number of the VXLAN packet. For all the inner flows of a particular VXLAN packet, first filter by VXLAN packets on the **App L4 Port** window to display all VXLAN packets. Identify the VXLAN packet you are interested in from the **Flows by Time** window. Expand the row, note the packet's VNI number, then remove the VXLAN filter and

---

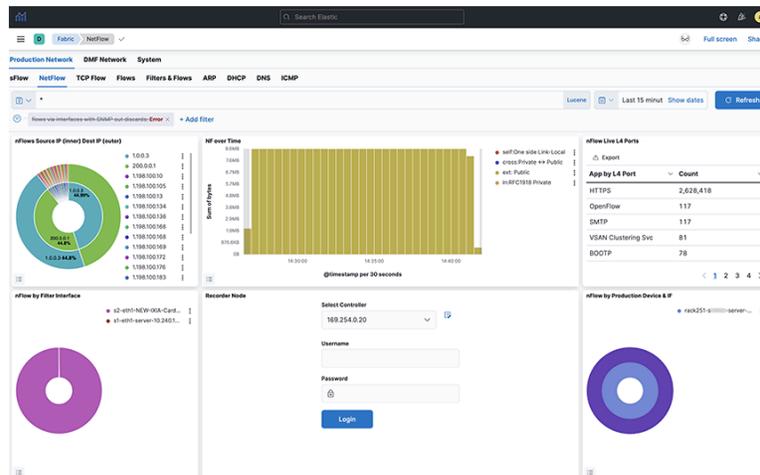
\* sFlow® is a registered trademark of Inmon Corp.

filter based on the VNI number. It will show both the outer flow of the VXLAN packet and all the inner flows associated with that VXLAN packet.

## 2.2 NetFlow and IPFIX

The system displays the following dashboard by clicking **NetFlow**:

**Figure 2-1: Production Network > NetFlow Dashboard**



Configure the NetFlow collector interface on the Arista Analytics Node to obtain NetFlow packets, as described in the [Setting up the NetFlow Collector on the Analytics Node](#) section.

The NetFlow dashboard summarizes information from the NetFlow messages sent to the Arista Analytics Node from the DANZ Monitoring Fabric controller or other NetFlow flow exporter and provides the following panels:

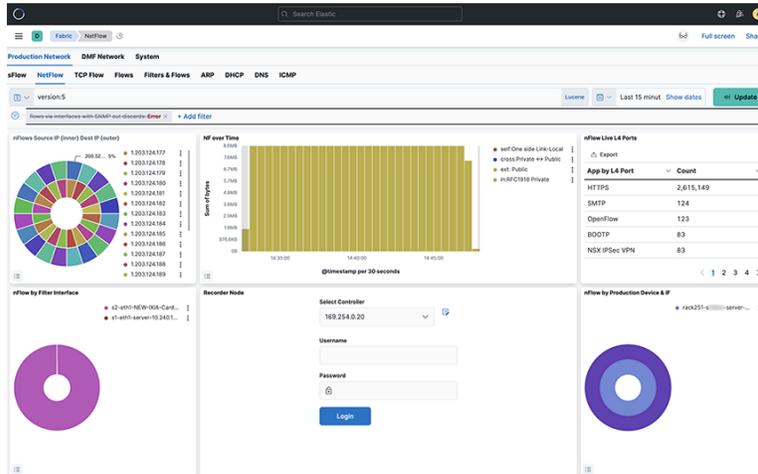
- nFlow Source IP (inner) Destination IP (outer)
- NF over Time
- nFlow Live L4 Ports
- nFlow by Filter Interface
- nFlow by Production Device & IF
- NF by QoS PHB
- NF by DPI App Name
- NF Top Talkers by Flow
- NF Detail



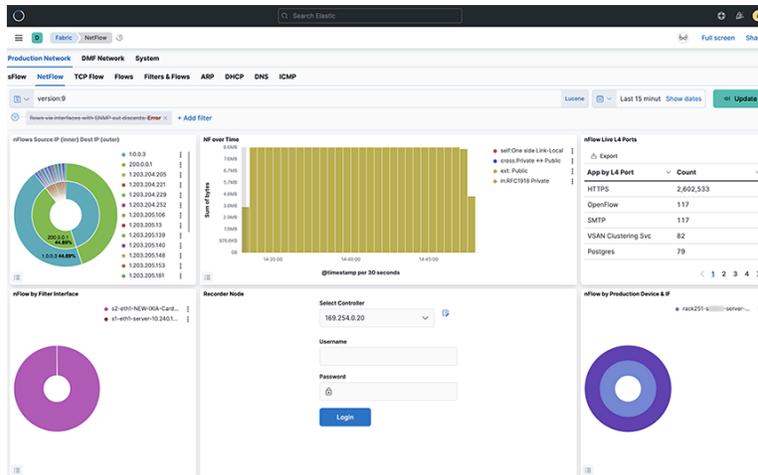
**Note:** To display the fields in the **nFlow by Filter Interface** panel for NetFlow V5 and IPFIX generated by the DMF Service Node appliance, **records-per-interface**, and **records-per-dmf-interface** knobs must be configured in the DANZ Monitoring Fabric controller.

Starting from the **BMF-7.2.1** release, the Arista Analytics Node can also handle NetFlow V5/V9 and IPFIX traffic. All of the flows represent a Netflow index. From the NetFlow Dashboard, filter rules apply to display specific flow information.

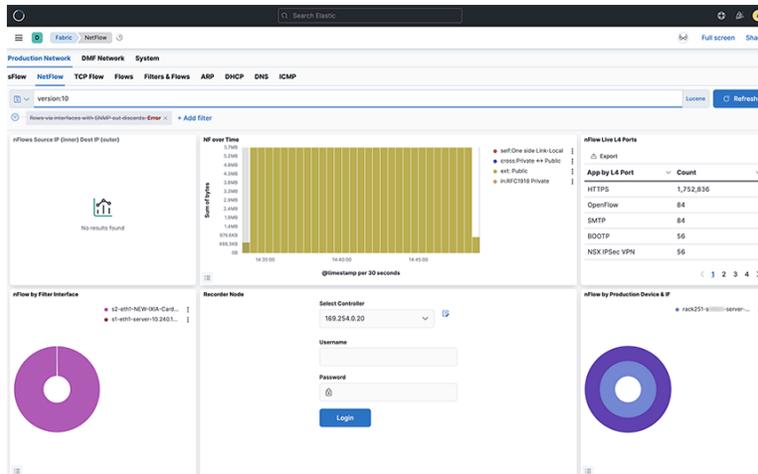
**Figure 2-2: NetFlow Version 5**



**Figure 2-3: NetFlow Version 9**



**Figure 2-4: NetFlow Version 10**



**Note:**

1. The Arista Analytics Node cluster listens to NetFlow v9 and IPFIX traffic on UDP **port 4739**. NetFlow v5 traffic learn on UDP **port 2055**.
2. Refer to **DANZ Monitoring Fabric 8.4 User Guide** for NetFlow and IPFIX service configuration.
3. Starting from the **DMF-8.1.0** release, Analytics Node capability augments in support of the following Arista Enterprise-Specific Information Element IDs:
  - 1036 -AristaBscanExportReason
  - 1038 -AristaBscanTsFlowStart
  - 1039 -AristaBscanTsFlowEnd
  - 1040 -AristaBscanTsNewLearn
  - 1042 -AristaBscanTagControl
  - 1043 -AristaBscanFlowGroupId

## 2.2.1 Consolidating Netflow V9/IPFIX records

You can consolidate NetFlow V9 and IPFIX records by grouping those with similar identifying characteristics within a configurable time window. This process reduces the number of documents published in Elasticsearch, decreases disk usage, and improves efficiency. This is particularly beneficial for long flows, where consolidations as high as 40:1 have been observed. However, enabling consolidation is not recommended for environments with low packet flow rates, as it may cause delays in the publication of documents.

The following configuration sets the load-balancing policy of Netflow/IPFIX traffic among nodes in DMF Analytics.

```
cluster:analytics# config
analytics(config)# analytics-service netflow-v9-ipfix
analytics(config-controller-service)# load-balancing policy source-hashing
```

The two settings are:

- **Source hashing:** forwards packets to nodes statistically assigned by a hashtable of their source IP address. Consolidation operations are performed on each node independently in source hashing.
- **Round-robin:** distributes the packets equally between the nodes if source-hashing results in significantly unbalanced traffic distribution. Round-robin is the default behavior.



**Note:** Configure the round-robin to lighten the load on the leader node when flow rate is higher than 10k/sec in cluster setup.

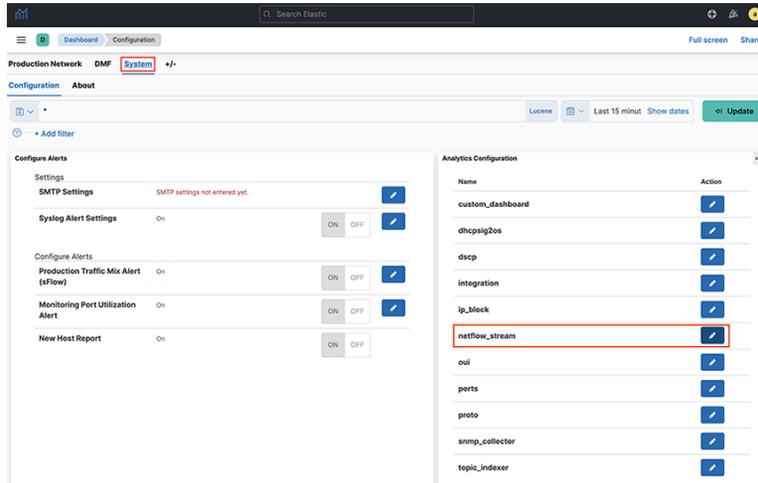


**Note:** This configuration doesn't apply to single-node deployments.

## Kibana Setup

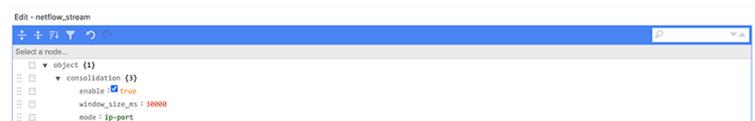
To perform the Kibana configuration, select the **System** > **Configuration** tab on the Fabric page and open the **Analytics Configuration** > **netflow\_stream** panel:

Figure 2-5: Kibana setup



For editing the netflow stream, go to the following tab:

Figure 2-6: Edit the netflow stream



There are three required settings:

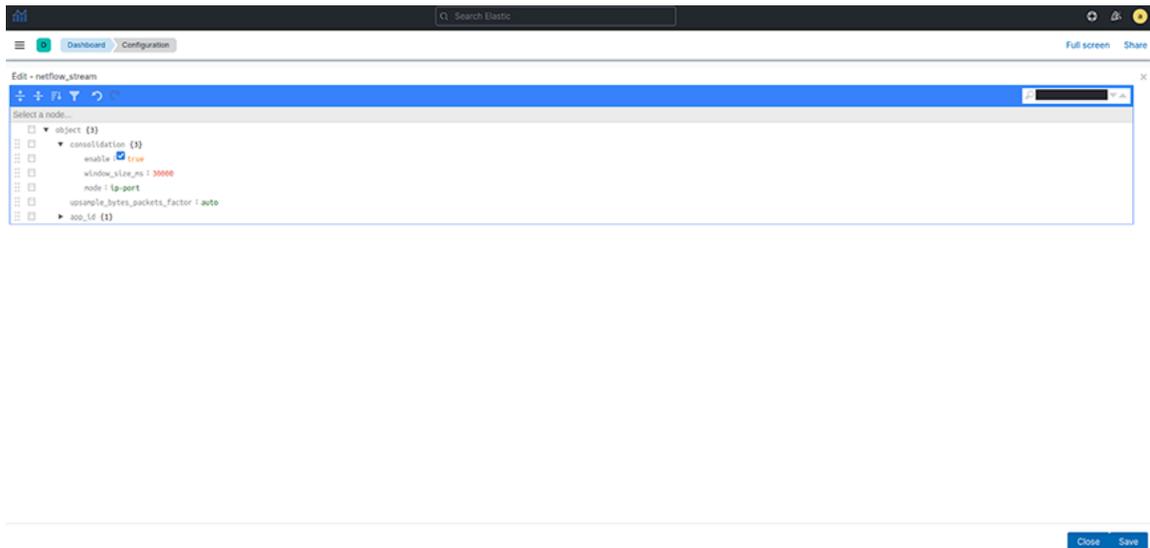
- **enable**: turn consolidation on or off.
- **window\_size\_ms**: adjust window size using the rate of Netflow V9/IPFIX packets per second the analytics node receives. The default window size is 30 seconds but measured in milliseconds.
- **mode**: There are three supported modes:
  - **ip-port**: records with the same source IP address, destination IP address, and IP protocol number. It also consolidates the lower numerical value of the source or destination Layer 4 port number with others.
  - **dmf-ip-port-switch**: records from common DMF Filter switches that meet **ip-port** criteria.
  - **src-dst-mac**: records with the same source and destination MAC addresses.



**Note:** It uses the mode when Netflow V9/IPFIX templates collect only Layer 2 fields.

Starting in **DMF-8.5.0**, the configuration mentioned above is set under a “**consolidation JSON**” object as follows:

**Figure 2-7: Consolidating Netflow**



### Consolidation Troubleshooting

If consolidation is enabled but does not occur, Arista Networks recommends creating a support bundle and contacting Arista TAC.

### Load-balancing Troubleshooting

If there are any issues related to load-balancing, Arista Networks recommends creating a support bundle and contacting Arista TAC.

## 2.2.2 NetFlow and IPFIX Flow with Application Information

This feature of Arista Analytics combines Netflow and IPFIX records containing application information with Netflow and IPFIX records containing flow information.

This feature improves the data visibility per application by correlating flow records with applications identified by the flow exporter.

This release supports only applications exported from Arista Networks Service Nodes. In a multi-node cluster, you must configure load balancing in the Analytics Node CLI command.

### Configuration

In a multi-node Analytics cluster, set the load-balancing policy of Netflow/IPFIX traffic to **source-hashing** as the **round-robin** policy may cause application information to be missing from the resulting flow documents in Elasticsearch.

```
analytics# config
analytics(config)# analytics-service netflow-v9-ipfix
analytics(config-an-service)# load-balancing policy source-hashing
```

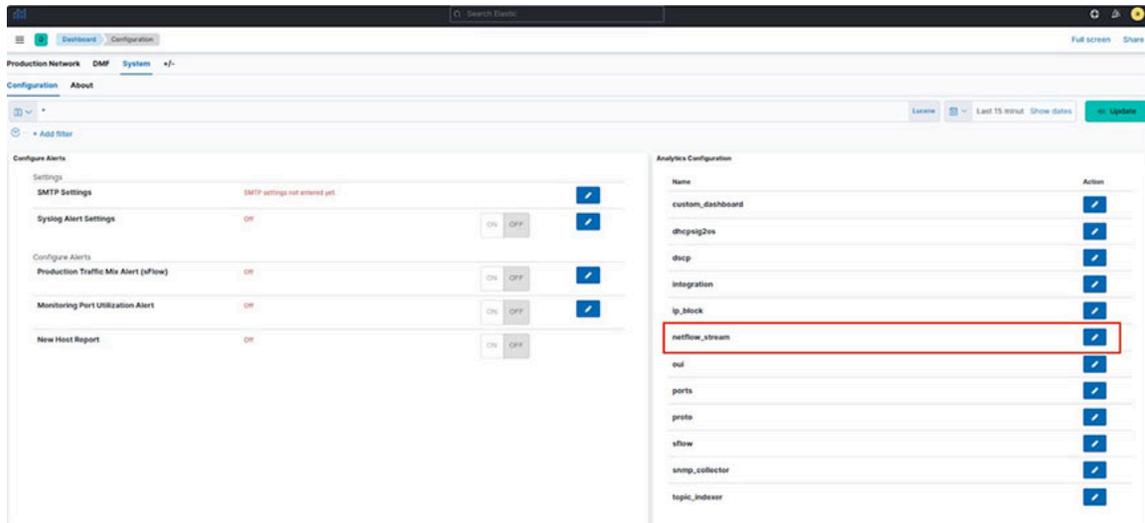


**Note:** This configuration doesn't apply to single-node deployments.

## Kibana Configuration

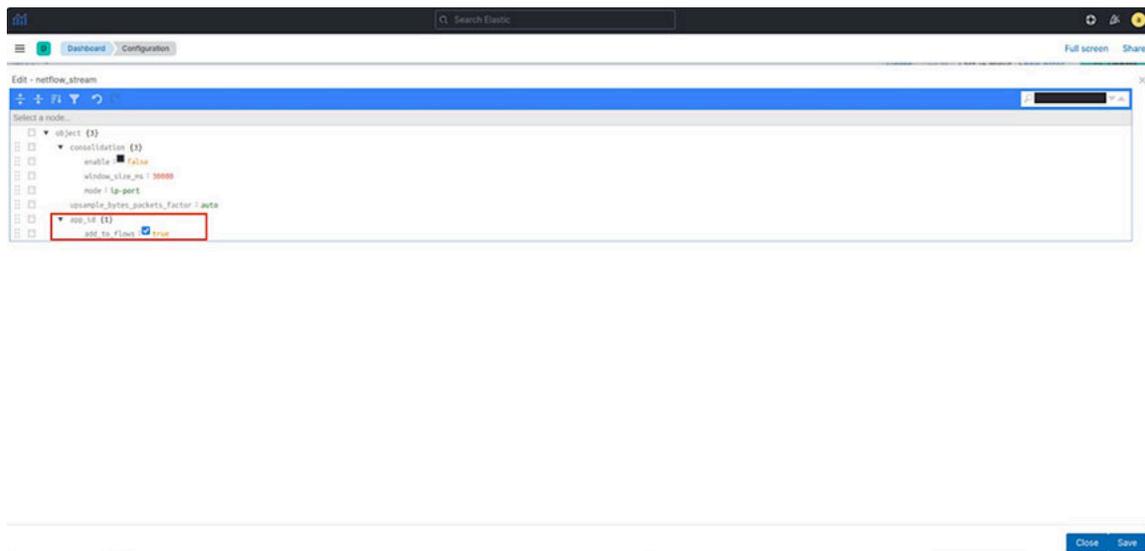
To perform the Kibana configuration, select the **System** > **Configuration** tab on the Fabric page and open the **Analytics Configuration** > **netflow\_stream** visualization.

**Figure 2-8: Dashboard - Netflow stream configuration**



Add the **app\_id** configuration object.

**Figure 2-9: Edit - Netflow stream**



In the **app\_id** configuration object, it requires the following setting:

- **add\_to\_flows**: Enables or turns off the merging feature.

## ElasticSearch Documents

Three fields display the application information in the final NetFlow/IPFIX document stored in ElasticSearch:

- **appScope**: Name of the NetFlow/IPFIX exporter.
- **appName**: Name of the application. This field is only populated if the exporter is NTOP.
- **appID**: Unique application identifier assigned by the exporter.

## Troubleshooting

If merging is enabled but does not occur, Arista Networks recommends creating a support bundle and contacting Arista TAC.

## Limitations

- Some flow records may not include the expected application information when configuring round-robin load balancing of Netflow/IPFIX traffic. Arista Networks recommends configuring the source-hashing load-balancing policy and sending all Netflow/IPFIX traffic to the Analytics Node from the same source IP address.
- Application information and flow records are correlated only if the application record is available before the flow record.
- Arista Networks only supports collecting application information from Netflow/IPFIX exporters: NTOP, Palo Alto Networks firewalls, and Arista Networks Service Node.
- This feature isn't compatible with the consolidation feature documented in the [Consolidating Netflow V9/IPFIX records](#). When merging with application information is enabled, consolidation must be disabled.

## 2.2.3 NetFlow and sFlow Traffic Volume Upsampling

Arista Analytics can upsample traffic volume sampled by NetFlow V9/IPFIX and sFlow. This feature provides better visibility of traffic volumes by approximating the number of bytes and packets from samples collected by the NetFlow V9/IPFIX or sFlow sampling protocols. It gives those approximation statistics along with the ElasticSearch statistics. The feature bases the approximations on the flow exporter's sampling rate or a user-provided fixed factor.



**Note:** When the rate of flow packets is low or for short flows, the approximations will be inaccurate.

The **DMF 8.5.0** release does not support the automated approximation of total bytes and packets for Netflow V9/IPFIX. If upsampling is needed, Arista Networks recommends configuring a fixed upsampling rate.

## NetFlow/IPFIX Configuration

To perform the Kibana configuration, select the **System** > **Configuration** tab on the Fabric page and open the **Analytics Configuration** > **netflow\_stream** visualization.

Figure 2-10: Dashboard - Netflow IPFIX configuration

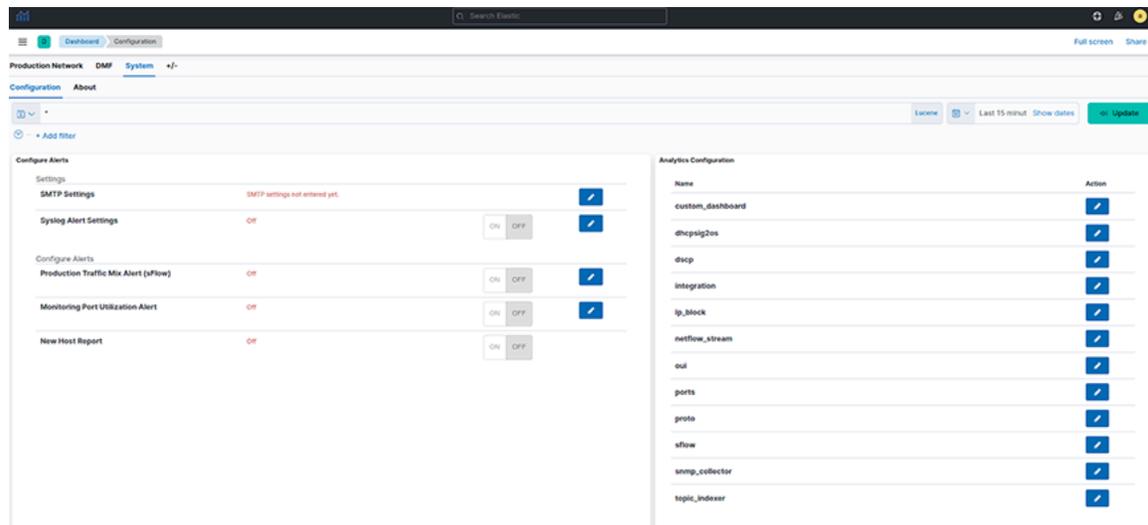
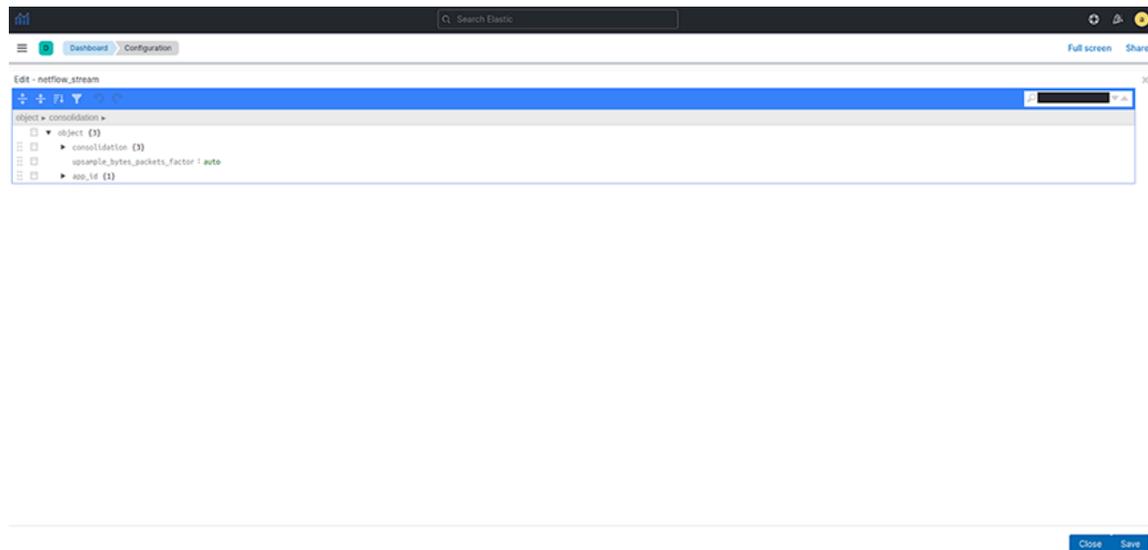


Figure 2-11: Edit - Netflow IPFIX



There is one required setting, **upsample\_byte\_packet\_factor**, with two possible options:

- **Auto**: This is the default option. **DMF 8.5.0** does not support automated upsampling for Netflow V9/IPFIX. Arista Networks recommends configuring an integer if upsampling is needed.
- **Integer**: Multiply the number of bytes and packets for each collected sample by this configured number.

## sFlow Configuration

To perform the Kibana configuration, select the **System** > **Configuration** tab on the Fabric page and open the **Analytics Configuration** > **sFlow** visualization.

Figure 2-12: Dashboard - sFlow configuration

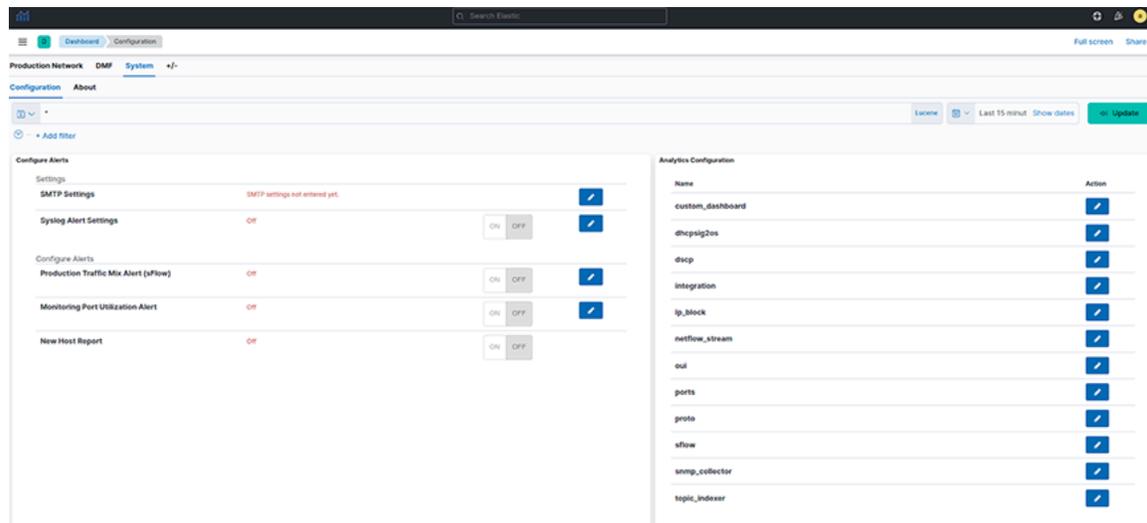
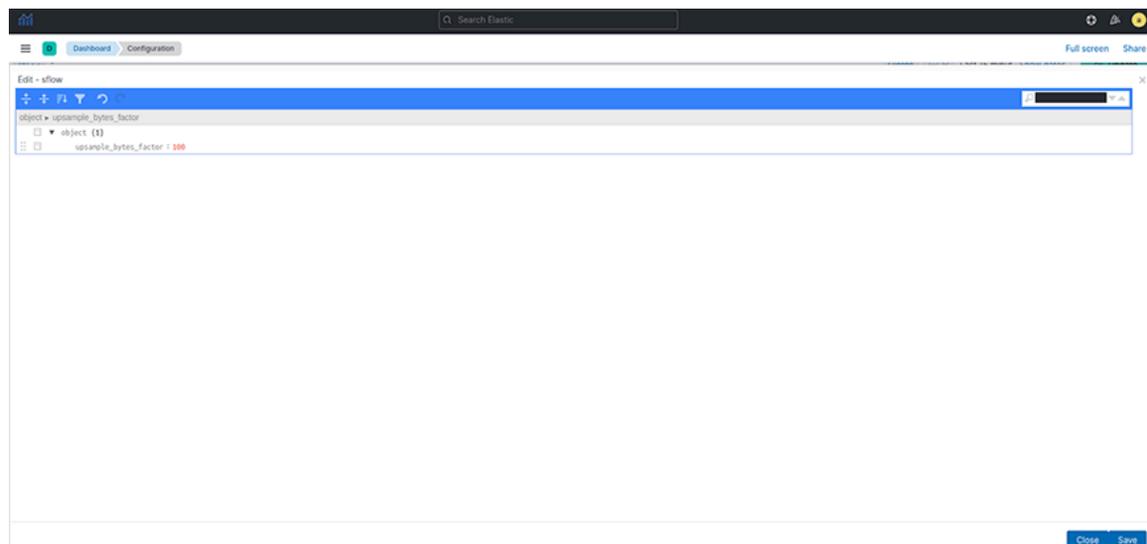


Figure 2-13: Edit - sFlow



There is one required setting, ***upsample\_byte\_packet\_factor***, with two possible options:

- **Auto**: Approximate the number of bytes and packets for each collected sample based on the collector's sampling rate. **Auto** is the default option.
- **Integer**: Multiply the number of bytes and packets for each collected sample by this configured number.

## Dashboards

### NetFlow Dashboard

The NetFlow dashboard is on the **Production Network** > **NetFlow** tab on the Fabric page. The following visualization will display upsampled statistics:

- **NF over Time**

- **NF Top Talkers by Flow**

**Figure 2-14: NF Detail visualization**

NF Detail								50 documents
Time	flow	proto	packets	upsampledPacketCount	bytes	upsampledByteCount	tos	
> Nov 21, 2023 @ 18:49:52.000	10.109.1.40>224.0.0.251	UDP	15	3,000	1.5KB	293KB		
> Nov 21, 2023 @ 18:49:52.000	10.109.1.40>224.0.0.251	UDP	15	3,000	1.5KB	293KB		
> Nov 21, 2023 @ 18:49:52.000	10.109.1.40>224.0.0.251	UDP	15	3,000	1.5KB	293KB		
> Nov 21, 2023 @ 18:49:52.000	10.109.1.40>224.0.0.251	UDP	15	3,000	1.5KB	293KB		
> Nov 21, 2023 @ 18:49:52.000	10.109.1.40>224.0.0.251	UDP	15	3,000	1.5KB	293KB		
> Nov 21, 2023 @ 18:49:52.000	10.109.1.40>224.0.0.251	UDP	15	3,000	1.5KB	293KB		
> Nov 21, 2023 @ 18:49:52.000	10.109.1.40>224.0.0.251	UDP	15	3,000	1.5KB	293KB		
> Nov 21, 2023 @ 18:49:52.000	10.109.1.40>224.0.0.251	UDP	15	3,000	1.5KB	293KB		

Rows per page: 50 < 1 of 1 >

The **DMF 8.5.0** release adds two new columns:

- **upsampledPacketCount**: Approximate total count of packets for a flow.
- **upsampledByteCount**: Approximate total count of bytes for a flow.



**Note:** In **DMF 8.5.0**, configuring upsampling to **Auto**, **upsampledByteCount**, and **upsampledPacketCount** will copy the **bytes** and **packets** column and display the values of **bytes** and **packets** in the graphs and tables of this dashboard.

### sFlow Dashboard

The sFlow dashboard is on the **Production Network > sFlow** tab on the Fabric page. The **Traffic over Time** visualization will display upsampled statistics.

**Figure 2-15: Flow by Time visualization**

Flows by Time							150 documents
Time	flow	sVendor	dVendor	bytes	upsampledByteCount	BTName	phb
> Nov 21, 2023 @ 19:07:34.418	145.132.153.226/KPN B.V.:50800>72.31.77.37/Charter Communications Inc:HTTPS			1,000	488.343	-	BE
> Nov 21, 2023 @ 19:07:31.418	93.66.226.4/Vodafone Italia S.p.A.:4499>89.231.137.231/MultiMedia Polska Sp. z o.o.:HTTPS			1.2KB	234.640	-	BE
> Nov 21, 2023 @ 19:07:28.419	152.194.85.78/Verizon Business:63580>86.214.156.18/Orange S.A.:HTTPS			1.5KB	166.563	-	BE
> Nov 21, 2023 @ 19:07:25.417	145.132.153.226/KPN B.V.:50800>72.31.77.37/Charter Communications Inc:HTTPS			1,000	488.343	-	BE
> Nov 21, 2023 @ 19:07:22.417	93.66.226.4/Vodafone Italia S.p.A.:4499>89.231.137.231/MultiMedia Polska Sp. z o.o.:HTTPS			1.2KB	234.640	-	BE
> Nov 21, 2023 @ 19:07:19.418	152.194.85.78/Verizon Business:63580>86.214.156.18/Orange S.A.:HTTPS			1.5KB	166.563	-	BE
> Nov 21, 2023 @ 19:07:16.417	145.132.153.226/KPN B.V.:50800>72.31.77.37/Charter Communications Inc:HTTPS			1,000	488.343	-	BE
> Nov 21, 2023 @ 19:07:13.416	93.66.226.4/Vodafone Italia S.p.A.:4499>89.231.137.231/MultiMedia Polska Sp. z o.o.:HTTPS			1.2KB	234.640	-	BE
> Nov 21, 2023 @ 19:07:10.415	152.194.85.78/Verizon Business:63580>86.214.156.18/Orange S.A.:HTTPS			1.5KB	166.563	-	BE
> Nov 21, 2023 @ 19:07:07.416	145.132.153.226/KPN B.V.:50800>72.31.77.37/Charter Communications Inc:HTTPS			1,000	488.343	-	BE

Rows per page: 50 < 1 of 3 >

The newly added **upsampledByteCount** represents a flow's approximate total count of bytes.

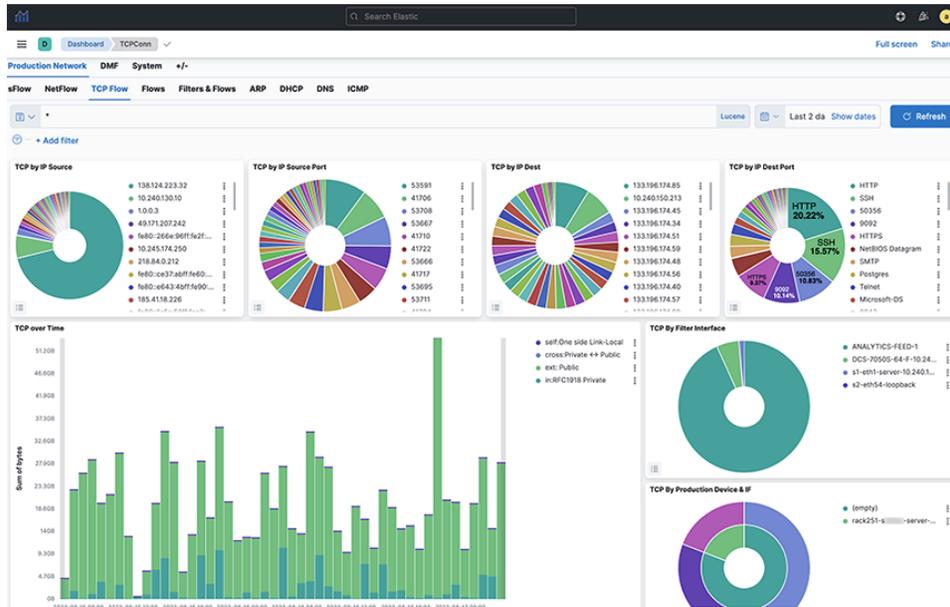
### Troubleshooting

Arista Networks recommends creating a support bundle and contacting Arista Networks TAC if upsampling isn't working correctly.

## 2.3 TCPFlow

Click the **TCPFlow** tab to display the following dashboard.

**Figure 2-16: Production Network > TCPFlow Dashboard**

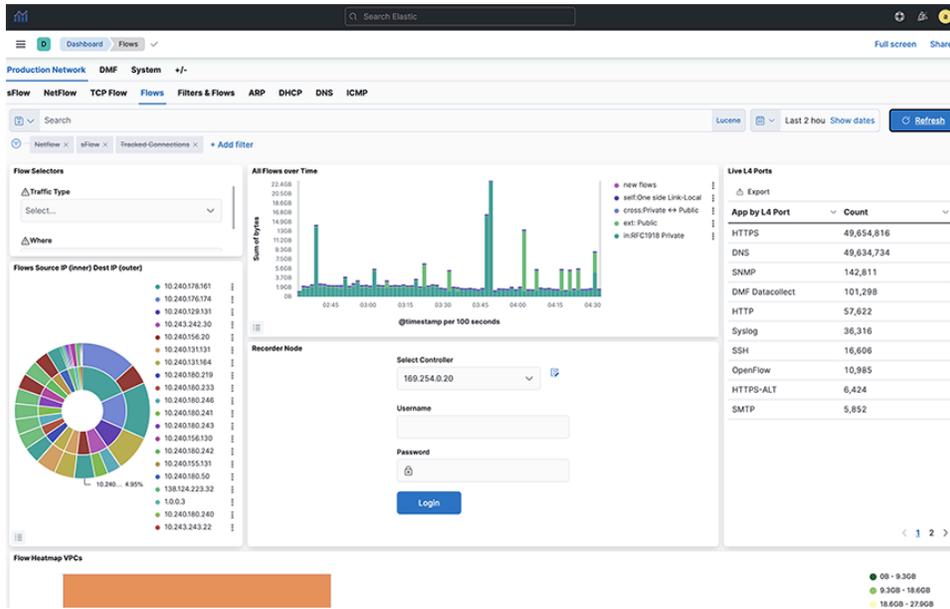


The information on the TCPFlow dashboard depends on TCP handshake signals and deduplicates. The **Filter Interface** visualization indicates the filter switch port where data is received. The switch description is specified in the Description attribute of each switch, configured on the DANZ Monitoring Fabric controller. **Device & IF** on this dashboard refers to the end device and depends on LLDP packets received.

## 2.4 Flows

Click the **Flows** tab to display the following dashboard.

**Figure 2-17: Production Network > Flows Dashboard**



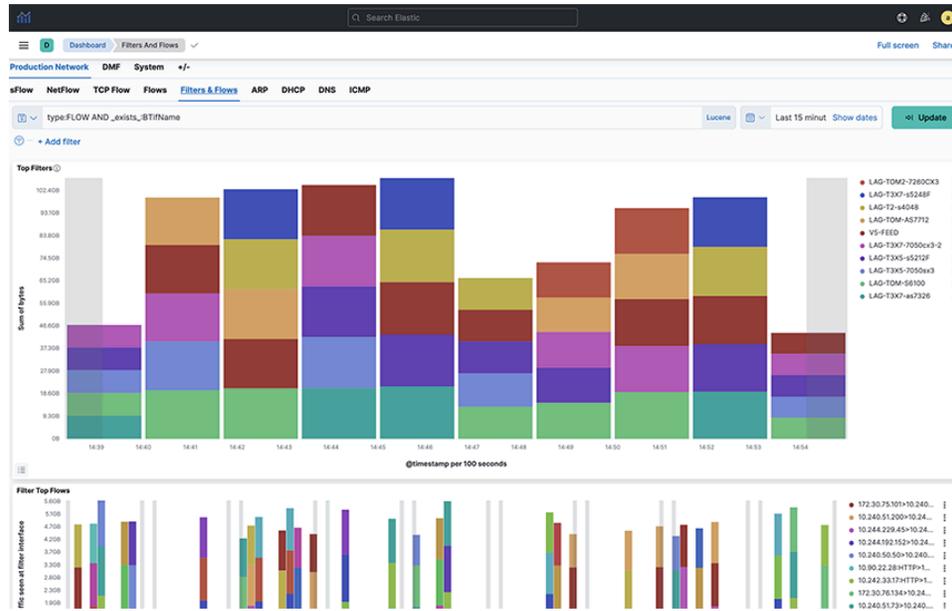
The **Flows Dashboard** summarizes information from sFlow and NetFlow messages and provides the following panels:

- All Flows Type
- All Flows Overtime
- All Flows Details

## 2.5 Filters & Flows

Click the **Filters & Flows** tab to display the following dashboard.

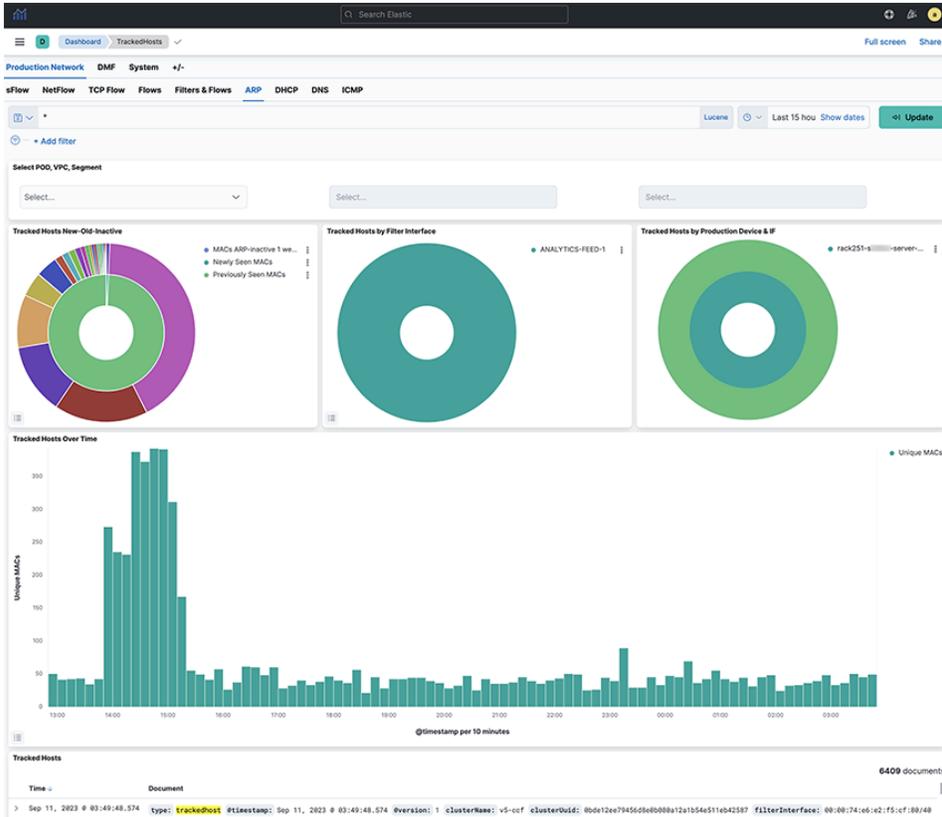
**Figure 2-18: Production Network > Filters & Flows Dashboard**



## 2.6 ARP

Click the **ARP** tab to display the following dashboard. This data correlates with the tracked host feature on the DANZ Monitoring Fabric controller. It shows all ARP data when you switch interface and production devices over time.

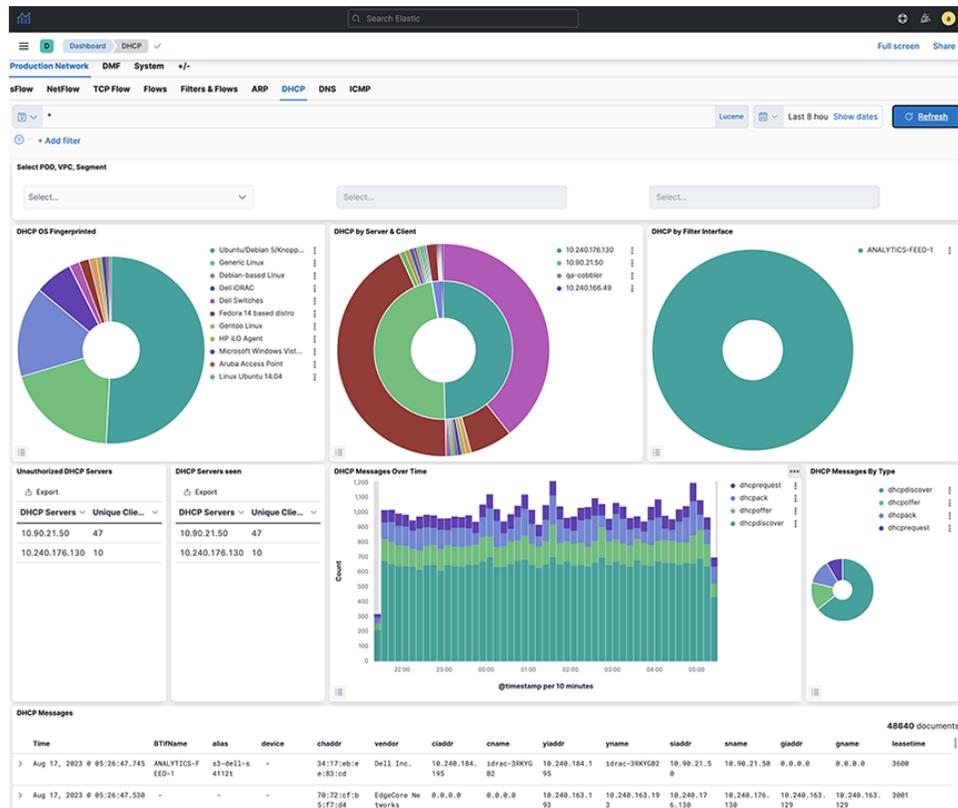
**Figure 2-19: Production Network > ARP Dashboard**



## 2.7 DHCP

Click the **DHCP** tab to display the following dashboard.

**Figure 2-20: Production Network > DHCP Dashboard**



**Note:** Operating systems on the network and data by filter interface and production device information are available.

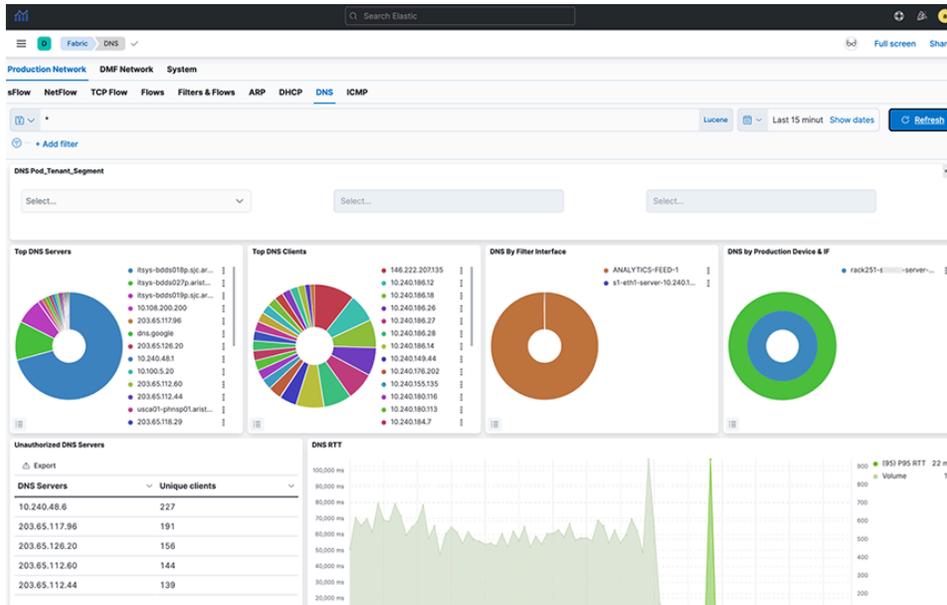
The **DHCP Dashboard** summarizes information from analyzing DHCP activity and provides the following panels:

- DHCP OS Fingerprinted
- DHCP Messages by Filter Interface
- DHCP Messages by Production Switch
- Non-whitelist DHCP Servers
- DHCP Messages Over Time
- DHCP Messages by Type
- DHCP Messages

## 2.8 DNS

Click the **DNS** tab to display the following dashboard.

**Figure 2-21: Production Network > DNS Dashboard**



The **DNS Dashboard** summarizes information from analyzing DNS activity and provides the following panels:

- DNS Top Servers
- DNS Top Clients
- DNS By Filter Interface
- DNS by Production Device & IF
- DNS Messages Over Time
- Unauthorized DNS Servers
- DNS RTT
- DNS All Messages
- DNS RCode Distro
- DNS QType Description
- DNS Top QNames

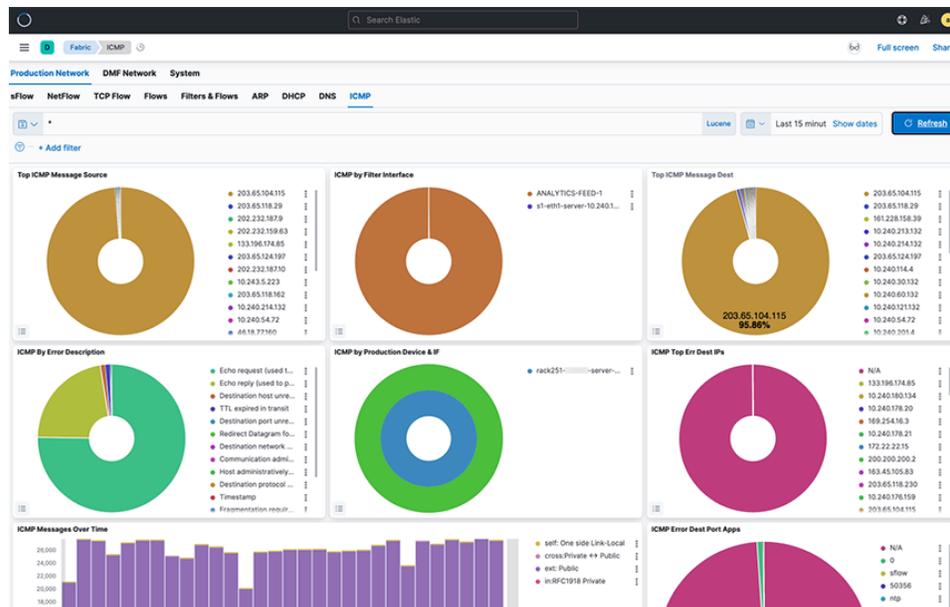


**Note:** The query and response packet timestamps compute the DNS RTT value. If a query packet does not answer by a response packet within **180** seconds, then the RTT value is set to **-1**.

## 2.9 ICMP

Click the **ICMP** tab to display the following dashboard.

**Figure 2-22: Production Network > ICMP Dashboard**



The **ICMP Dashboard** summarizes information from analyzing ICMP activity and provides the following panels:

- Top ICMP Message Source
- ICMP by Filter Interface
- Top ICM Message Dest
- ICMP by Error Description
- ICMP by Production Switch
- ICMP Top Err Dest IPs
- ICMP Top Err Dest Port Apps
- ICMP Messages Over Time
- ICMP Table

# Using the DMF Recorder Node with Analytics

---

This chapter describes Arista Analytics to use with the DANZ Monitoring Fabric Recorder Node. It includes the following sections.

- [Overview](#)
- [General Operation](#)
- [Using Recorder with Analytics](#)

## 3.1 Overview

The DMF Recorder Node records packets from the network to disk and recalls specific from disk quickly, efficiently, and at scale. A single DANZ Monitoring Fabric controller can manage multiple DMF Recorder Nodes, delivering packets for recording through DANZ Monitoring Fabric policies. The controller also provides central APIs for interacting with DMF Recorder Nodes to perform packet queries across one or multiple recorders and for viewing errors, warnings, statistics, and the status of connected recorder nodes.

A DANZ Monitoring Fabric policy directs matching packets to one or more recorder interfaces. The DMF Recorder Node interface defines the switch and port used to attach the recorder to the fabric. A DANZ Monitoring Fabric policy treats these as delivery interfaces.

Both NetFlow and TCPflow dashboards have the recorder node visualization.

## 3.2 General Operation

To retrieve packets from the DMF Recorder Node for analysis using Arista Analytics, select the Controller and log in from the **Recorder Node** window on the **NetFlow** or **Flows** dashboard. To add a new Controller,

click the small **Select Controller** icon and add the Controller. After logging in to the DMF Recorder Node, the system displays the following dialog:

**Figure 3-1: DMF Recorder Node**

The screenshot shows a web interface for configuring a query in the DMF Recorder Node. At the top left, there is a '+ Add filter' link. Below it, a 'Back to Login' link is visible. The interface displays the 'Oldest Packet' as 'Jul 7th 2023, 12:31PM' and the 'Latest Packet' as 'Aug 17th 2023, 09:29AM'. A horizontal menu contains tabs for 'Size', 'AppID', 'Packet Data', 'Packet Objects', 'Replay', and 'Flow Analysis', with 'Size' currently selected. Below the menu is a date range selector set to 'Last 15 minutes' with a 'Show dates' link and a 'Refresh' button. There are two input fields: 'IP Protocol #' and 'Community ID'. Below these are 'Source Info >' and '< Destination Info' buttons, a 'Bi-directional' toggle switch which is checked, and an 'Additional Parameters' dropdown menu. At the bottom, there are three buttons: 'Abort' (red), 'Clear' (light blue), and 'Submit' (dark blue). Below the buttons, there is a link '> Query Preview: Size'.

The **Recorder Node** window can compose and submit a query to the DMF Recorder Node. Use any of the fields shown to create a query and click **Submit**. The **Switch Controller** link at the bottom of the dialog can log in to a different DMF Recorder Node.

Use the **Recorder Summary** query to determine the number of packets in the recorder database. Then, apply filters to retrieve a reasonable number of packets with the most interesting information.

You can modify the filters in the recorder query until a **Size** query returns the most beneficial number of packets.

### Query Parameters

The following parameters are available for queries:

- **Query Type**
  - **Size:** Retrieve a summary of the matching packets based on the contents and search criteria stored in the recorder node. Here, Size refers to the total frame size of the packet.
  - **AppID:** Retrieve details about the matching packets based on the contents and search query in the recorder node datastore, where the packets are stored. Use this query to see what applications are in encrypted packets.
  - **Packet Data:** Retrieve the raw packets that match the query. At the end of a search query, it generates a URL pointing to the location of the pcap if the search query is successful.
  - **Packet Objects:** Retrieve the packet objects that match the query. At the end of a search query, it generates a URL pointing to the location of the objects (images) if the search query is successful.
  - **Replay:** Identify the Delivery interface in the field that appears, where the replayed packets are forwarded.
  - **FlowAnalysis:** Select the flow analysis type (HTTP, HTTP Request, DNS, Hosts, IPv4, IPv6, TCP, TCP Flow Health, UDP, RTP Streams, SIP Correlate, SIP Health).
- **Time/Date Format:** Identify the matching packets' time range as an absolute value or relative to a specific time, including the present.
- **Source Info:** Match a specific source IP address / MAC Address / CIDR address.

- 
- **Bi-directional:** Enabling this will query bi-directional traffic.
  - **Destination Info:** Match a specific destination IP address / MAC Address / CIDR address.
  - **IP Protocol:** Match the selected IP protocol.
  - **Community ID:** Flow hashing.

#### Additional Parameters

- **VLAN:** Match the VLAN ID.
- **Outer VLAN:** Match the outer VLAN ID when multiple VLAN IDs exist.
- **Inner/Middle VLAN:** Match the inner VLAN ID of two VLAN IDs or the middle VLAN ID of three VLAN IDs.
- **Innermost VLAN:** Match innermost VLAN ID of three VLAN IDs.
- **Filter Interfaces:** Match packets received at the specified DANZ Monitoring Fabric filter interfaces.
- **Policy Names:** Match packets selected by the specified DANZ Monitoring Fabric policies.
- **Max Size:** Set the maximum size of the query results in bytes.
- **Max Packets:** Limits the number of packets the query returns to this set value.
- **MetaWatch Device ID:** Matches on device ID / serial number found in the trailer of the packet stamped by the MetaWatch Switch.
- **MetaWatch Port ID:** Matches on application port ID found in the trailer of the packet stamped by the MetaWatch Switch.
- **Packet Recorders:** Query a particular DMF Recorder Node. Default is none or not selected; all packet recorders configured on the DANZ Monitoring Fabric receive the query.
- **Dedup:** Enable/Disable Dedup.
- **Query Preview:** After expanding, this section provides the Stenographer syntax used in the selected query. You can cut and paste the Stenographer query and include it in a REST API request to the DMF Recorder Node.

### 3.3 Using Recorder with Analytics

For interactive analysis, any set of packets exceeding **1 GB** becomes unwieldy. To reduce the number of packets to a manageable size, complete the following steps:

1. Use the Summary query to determine the number of packets captured by the Recorder. Apply filters until the packet set is manageable (less than **1 GB**).
2. Search over the metadata from all sources and analyze it to retrieve a limited and useful set of packets based on source address, destination address, timeframe, and other filtering attributes.
3. Submit the Stenographer query, which is used by the DMF Recorder Node and automatically composed by Arista Analytics.

You can perform flow analysis without downloading the packets from Recorder. Select specific rows to show Throughput, RTT, Out of order, and Re-transmissions. Packet varieties like HTTP, HTTP request, DNS, Hosts, IPv4, IPv6, TCP, TCPFlow Health, UDP, RTP Streams, SIP Correlate, and SIP Streams analyze the flows. Then, sort and search as required and save to CSV for later analysis. You can search over a given duration of time for the IP address by exact match or prefix match.

Replay set direct large packets to an archive for later analysis; this frees up the Recorder to capture a new packet set.

Use DMF Recorder Node to identify the applications on your network that are encrypting packets. Use a Recorder Detail query to see the applications with encrypted packets.

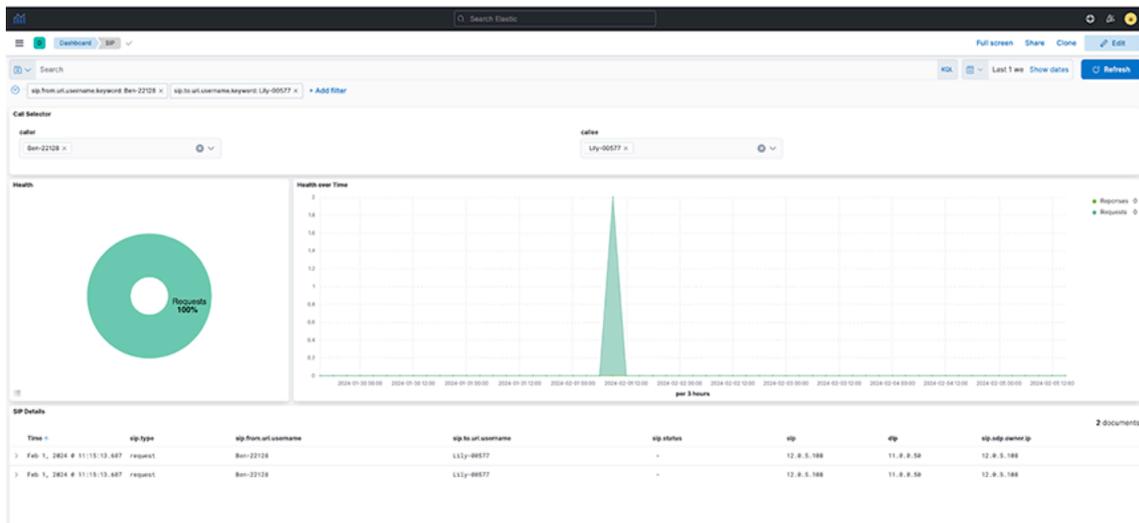
Refer to the ***DANZ Monitoring Fabric Deployment Guide*** for information about installing and setting up the DMF Recorder Node. For details about using the Recorder from the DANZ Monitoring Fabric Controller GUI or CLI, refer to the ***DANZ Monitoring Fabric User Guide***.

### 3.4 Analyzing SIP and RTP for DMF Analytics

This feature describes how Session Initiation Protocol (SIP) packets are parsed in a DANZ Monitoring Fabric (DMF) Analytics Node deployment and presented in a dashboard to allow the retrieval of data packets conveying voice traffic (RTP) from the DMF Recorder Node (RN). DMF accomplishes this by showing logical call information such as the call ID, phone number, and username. After retrieving the SIP record, the associated IP addresses are used to retrieve packets from the RN and opened in Wireshark for analysis.

Kibana has the **SIP** dashboard.

**Figure 3-2: SIP Dashboard**



#### DMF Preconditions

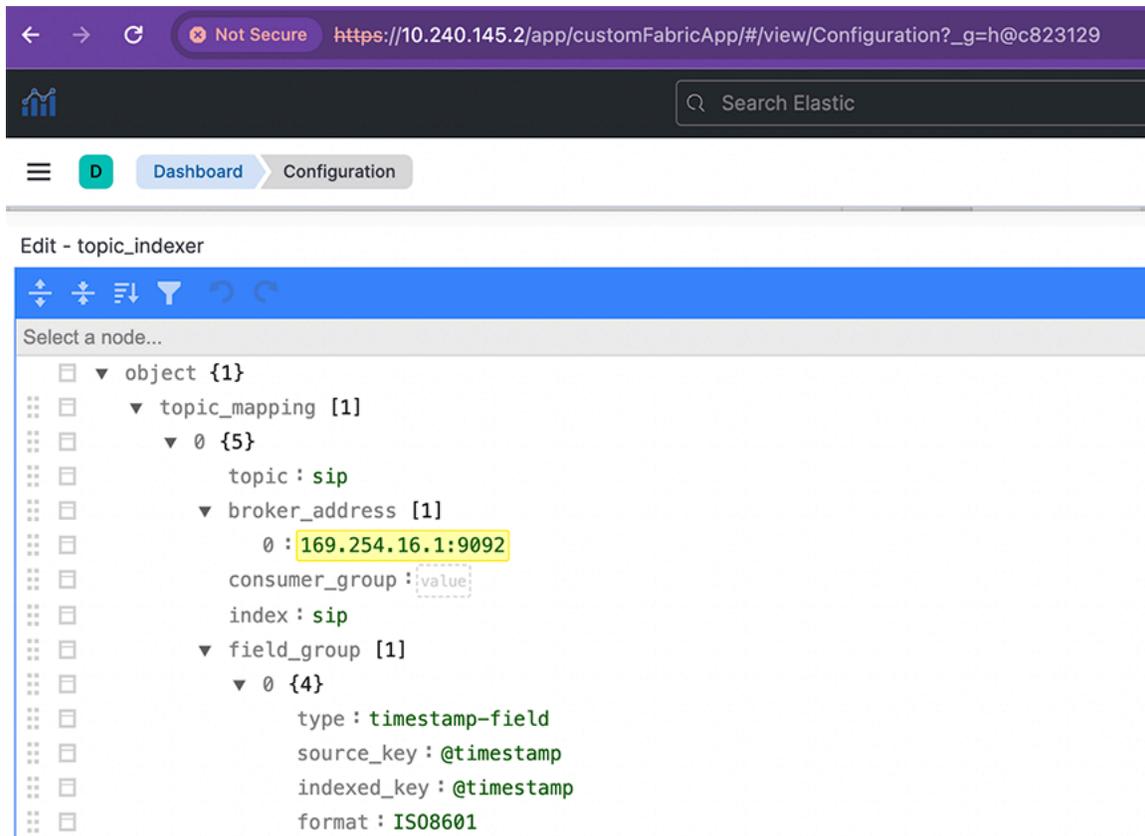
The feature requires a physical connection from the DMF Delivery Switch to the 10G Analytics Node (AN) Collector interface.

- Policy configured to filter for SIP traffic (UDP port 5060) such that low-rate traffic (< 1Gbps) is delivered to AN via collector interface with a filter on the Layer 4 port number or UDF.
- LAG will send SIP Control Packets to 1, 3, and 5 AN Nodes with symmetric hashing enabled and without hot-spotting.
- Recorder Node to receive SIP and Control packets recorded with standard key fields.

## Configuration

Configure SIP using the **broker\_address**, **timestamp-field**, and **field\_group** to enable the feature. Refer to [Field Details](#) for more information on **broker\_address**.

Figure 3-3: Edit-topic indexer



## Limitations

The **AN DMF 8.5.0** release supports this feature.

- There is no toggle switch to turn this feature on or off.

## Managing the NetFlow Dashboard

---

This chapter manages NetFlow and provides an efficient way to use the NetFlow dashboard.. Arista Analytics acts as a NetFlow collector for any agent or generator configured with the Analytics server IP address as a collector. It includes the DMF Service Node and any third-party NetFlow agent. This chapter has the following sections:

- [NetFlow Optimization](#)
- [Viewing Filter Interface Information on the NetFlow Dashboard](#)
- [Displaying Flows with Out-Discards](#)

### 4.1 NetFlow Optimization

Arista Analytics may consolidate NetFlow records to improve performance.

The Analytics server/cluster consolidates flows received within two seconds into a single flow when the source and destination IP addresses are the same or the source or destination L4 protocol port is the same.

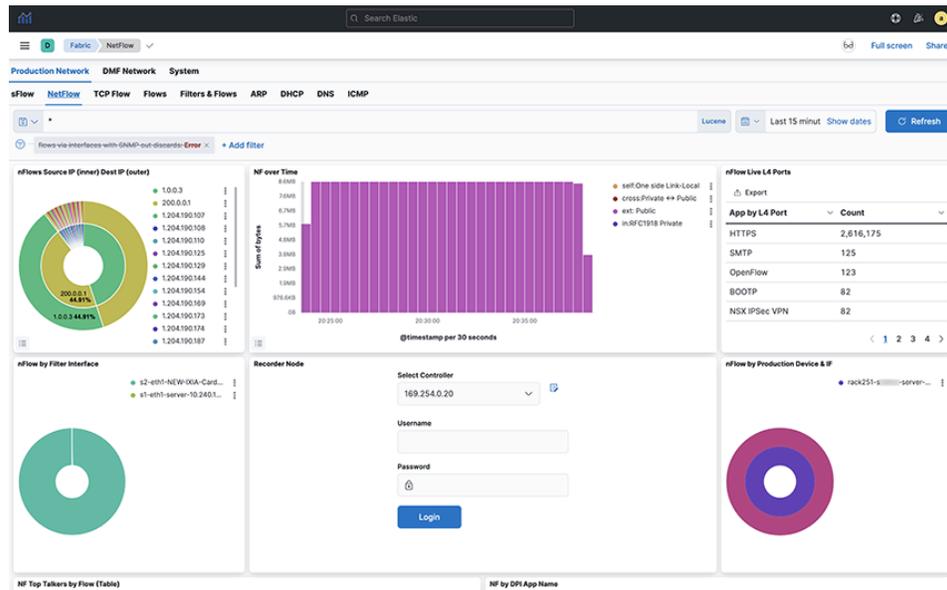
For example, ten flows received by the Analytics server within thirty second period are consolidated into a single flow if the source and destination IP addresses and destination port are the same for all the flows and only the source ports are different or if the source and destination IP addresses and source port are the same for all the flows and only the destination ports are different. This consolidated flow displays as a single row.

By default, the NetFlow Optimization is enabled for Netflow v5 and disabled for Netflow v9 and IPFIX. To allow the Netflow Optimization for Netflow v9 and IPFIX, refer to [Consolidating Netflow V9/IPFIX records](#) section.

This consolidation improves Analytics NetFlow performance, allowing more efficient indexing and searching of NetFlow information.

The following figure shows the **NF Detail** window on the **NetFlow** dashboard, which provides an example of NetFlow information with optimization.

**Figure 4-1: Analytics NetFlow Optimization**



## 4.2 Viewing Filter Interface Information on the NetFlow Dashboard

Add the filter interface name to the **NetFlow** dashboard to see hop-by-hop forwarding of flows for NetFlow traffic coming from the DMF Service Node for a specific flow. Arista Analytics then shows the filter interface name associated with that flow. If the flow goes through two hops, then two filter interface names are displayed for the flow.

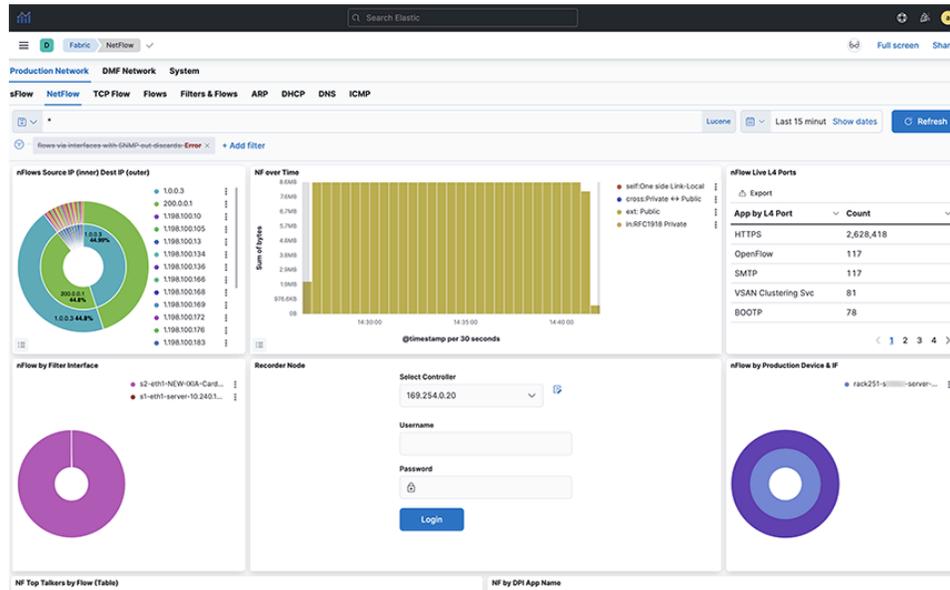
### 4.2.1 Displaying Filter Interface Names

The **nFlow by Filter Interface** window on the **NetFlow** dashboard, shown below, can display the filter interface name where traffic is coming in for the NetFlow service. To display this information, enable the records-per-interface option in the NetFlow managed service configuration on the DANZ Monitoring Fabric controller using the commands shown in the following example.

```
controller(config)# managed-service netflow-managed-service
controller(config-managed-srv)# service-action netflow netflow-delivery-int
```

```
controller(config-managed-srv-netflow)# collector 10.8.39.101 udp-port 2055 mtu
1500 records-per-interface
```

Figure 4-2: Production Network > NetFlow Dashboard with Filter Interface Name



### NetFlow Managed Service Records-per-interface Option

The following example displays the *running-config* for this configuration.

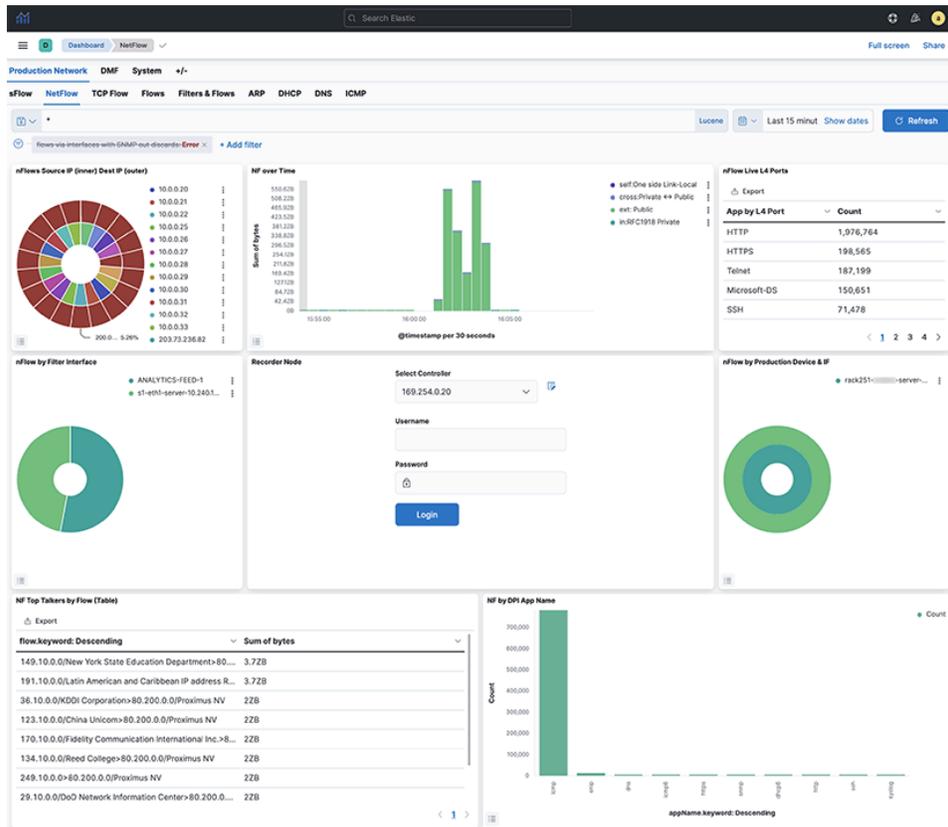
```
! managed-service
managed-service netflow-managed-service
service-interface switch 00:00:4c:76:25:f5:4b:80 ethernet4/3:4
!
service-action netflow netflow-delivery-int
collector 10.8.39.101 udp-port 2055 mtu 1500 records-per-interface
controller(config)# sh running-config bigtap policy netflow-policy
! policy
policy netflow-policy
action forward
filter-interface filter-int-eth5
use-managed-service netflow-managed-service sequence 1 use-service-delivery
1 match any
```

After enabling this option, the **nFlow by Filter Interface** window, shown above, displays the filter interface identified in the policy that uses the NetFlow managed service.

The production device port connected to the filter interface sends LLDP messages, Arista Analytics also displays the production switch name and the production interface name attached to the filter interface in the **nFlow by Production Switch & IF** window.

In the example below, **wan-tap-1** displays in the **nFlow by Filter Interface** window. The production device **N1524-WAN** and the interface **Gi1/0/1**, connected to filter interface **wan-tap-1**, are displayed in the **nFlow by Production Switch & IF** window.

**Figure 4-3: Production Network > NetFlow Dashboard with Filter Interface Name**



## 4.2.2 NetFlow Traffic Coming from Third-party Devices

This section displays third-party device and interface names. It shows hop-by-hop forwarding of flows when NetFlow traffic comes from a third-party device. For a query for a specific flow, Arista Analytics shows the device and interface names associated with that flow. If the flows go through two hops, it displays the device and interface names associated with flows.

Arista Analytics can act as a NetFlow collector for third-party devices. In this case, Arista Analytics displays third-party device management IP addresses and the interface index (iFindex) of the interface for each NetFlow-enabled third-party device.

For example, the **nFlow by Production Device & IF** window shows that **10.8.39.198** is the third-party device that forwards NetFlow traffic. The iFindex of the interface on that device where NetFlow is enabled is **0, 2, 3, 4**.

To discover the device name and the actual interface name rather than the iFindex, Arista Analytics automatically does an SNMP walk by getting the third-party device management IP from flow information. By default, Analytics uses the SNMP community name **public** to get the device name and interface name. If the SNMP community name of the third-party device is not **public**, change it in the Arista Analytics SNMP collector configuration.



**Note:** **AN DMF 8.3.0** release supports both SNMPv2 and SNMPv3.



**Note:** For IPFIX and nFlow v9, configure the third-party device to send the iFIndex. The Analytics node will do an SNMP walk to get the interface names associated with that iFIndex. By default, the iFIndex is not sent with IPFIX or nFlow v9. For example, to send the iFIndex for IPFIX and nFlow v9, enable `match interface input snmp` and `match interface output snmp` under **flow record** configuration on the third-party device.

#### DMF Analytic > System > Configuration > Analytic Configuration > snmp\_collector

Arista Analytics then performs SNMP polling and displays the third-party device name and the actual interface name in the **nflow by Production Device & IF** window.

To perform the SNMP configuration, complete the following steps:

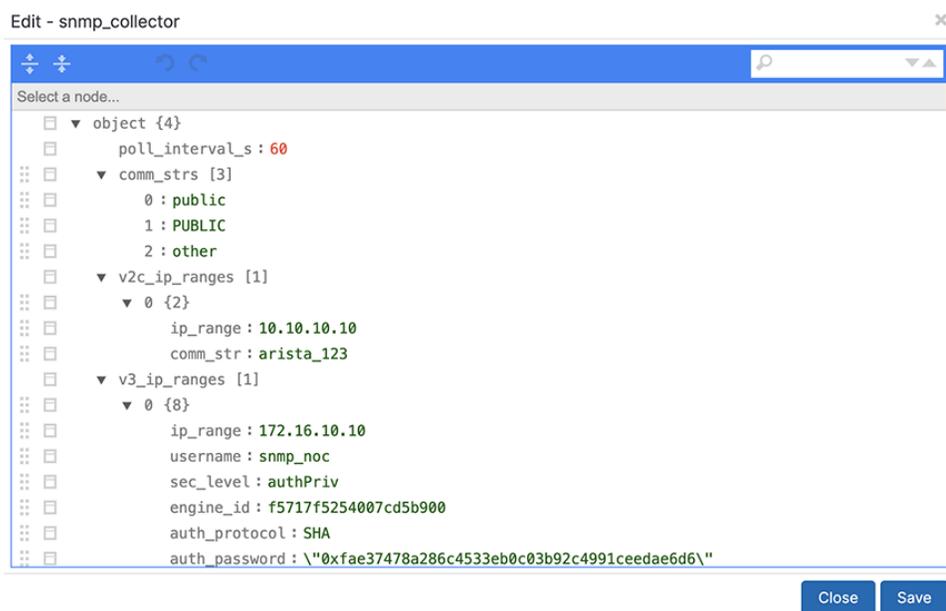
1. On the screen shown below, click **DMF Analytic > System > Configuration > Analytic Configuration > snmp\_collector > Edit**.

**Figure 4-4: Analytic snmp\_collector config**

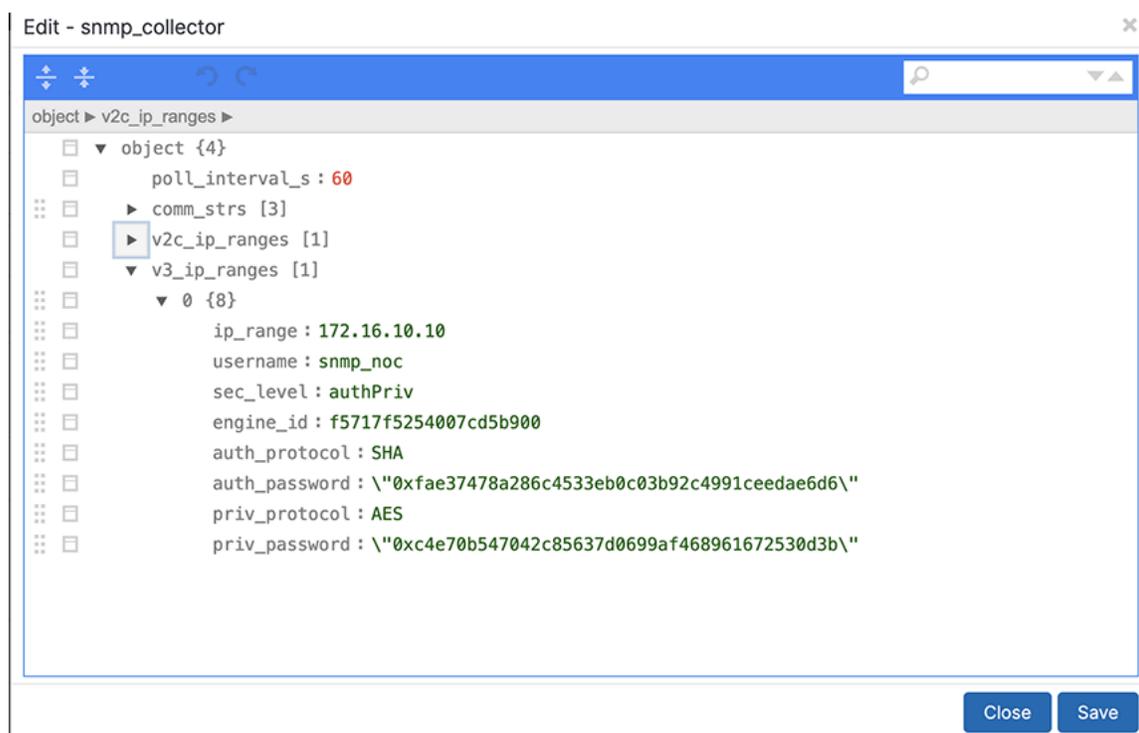
Analytics Configuration	
Name	Action
custom_dashboard	
dhcpsig2os	
dscp	
integration	
ip_block	
netflow_stream	
oui	
ports	
proto	
snmp_collector	
topic_indexer	

The system displays the following edit dialog.

**Figure 4-5: Analytic Configuration > snmp\_collector > Edit Dialog (SNMPv2 Configuration)**



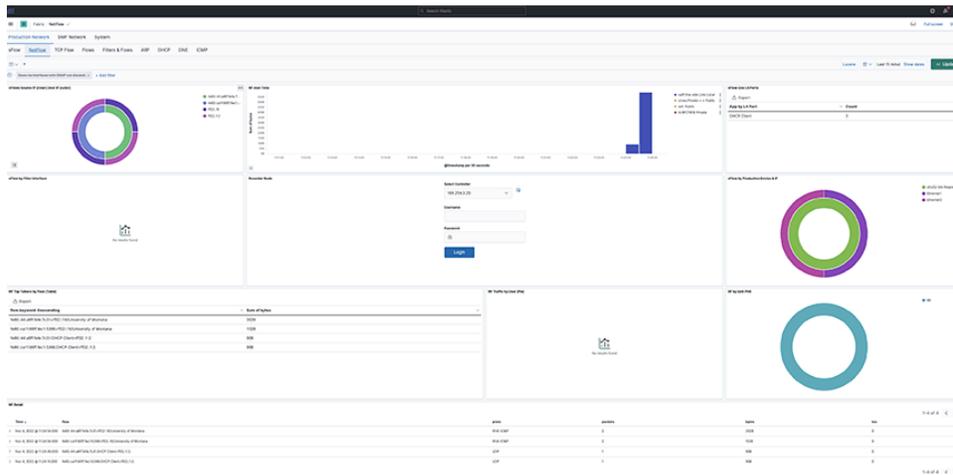
**Figure 4-6: Analytic Configuration > snmp\_collector > Edit Dialog (SNMPv3 Configuration)**



2. Click the community string **public** to change it to a different value as shown in the following dialog. By default, the SNMP collector polls devices every **60** seconds.
3. To change the SNMP poll interval, click the value **60**, change it to the preferred value, and click **Save**.

After completing this configuration, the third-party device is polled for the device name and interface name, **nflow by Production Device & IF** window displays it.

**Figure 4-7: Analytic Configuration > snmp\_collector > Edit Dialog**



### 4.3 Displaying Flows with Out-Discards

The **NetFlow** dashboard allows displaying flows with out-discards when the NetFlow packets come from third-party devices. To display this information, use the **flows via interfaces with SNMP out-discards** tab at the top of the Arista Analytics **NetFlow** dashboard.

To display the flows with out-discards, click the **flows via interfaces with SNMP out-discards** tab and click the **Re-enable** button. This window displays the flows with out-discards.

## Advanced Feature Dashboard

This chapter manages Latency and Drop Differ dashboard for DMF Analytics Node for keeping the records for NetFlow. This chapter has the following section:

- [Latency Differ and Drop Differ Dashboard](#)

### 5.1 Latency Differ and Drop Differ Dashboard

The DANZ Monitoring Fabric (DMF) Latency Differ Dashboard and Drop Differ Dashboard feature provides a near real-time visual representation of latency and drops in the DMF Analytics Node (AN) dedicated to NetFlow Records.

For a given flow, it reports the latency and drop of packets over time between two existing tap points (**A**, **B**), with network flows traversing the managed network from **A** towards **B**.

This feature introduces the concept of **DiffPair**, defined as a flow from **A** towards **B**.

The Dashboards provide clear, concise information about the flows. The data helps determine which applications are running slow and identifies peak times. A configurable mechanism alerts on abnormal drops and latency.

#### Introducing DiffPair

When identifying the flows between two tap points or filter interfaces, the aggregation occurs as **A** towards **B** pairs. It implies that a flow originating from point **A** will be received at point **B**. The term **DiffPair** is employed to visualize this flow as a cohesive set. This newly introduced field in the flow data selects the ingress and egress tap points encompassing a flow in between. The utilization of this **DiffPair** facilitates tap point filtering and comparison.



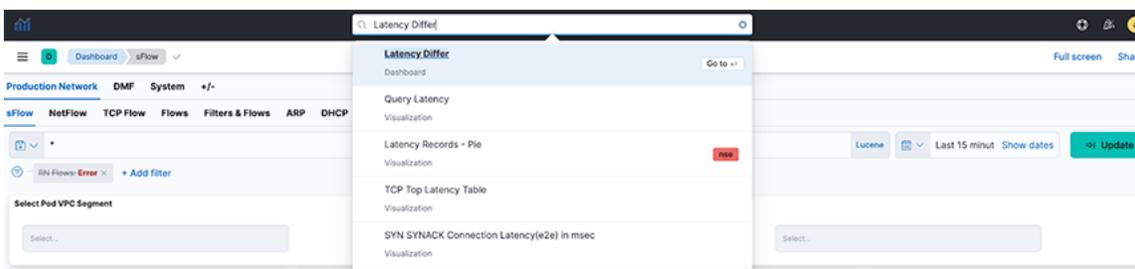
**Note:** It is important to verify the accuracy of the **DiffPair** data flowing between the tap points when comparing source data to the destination data.

#### Latency Differ Dashboard

Locate the **Latency Differ** dashboard by searching for the term **Latency Differ**.

The dashboard combines a visual representation of NetFlow Latency data in two views. The upper view displays individual flows, while the lower view aggregates A towards B pairs (**A > B**) or **DiffPair**.

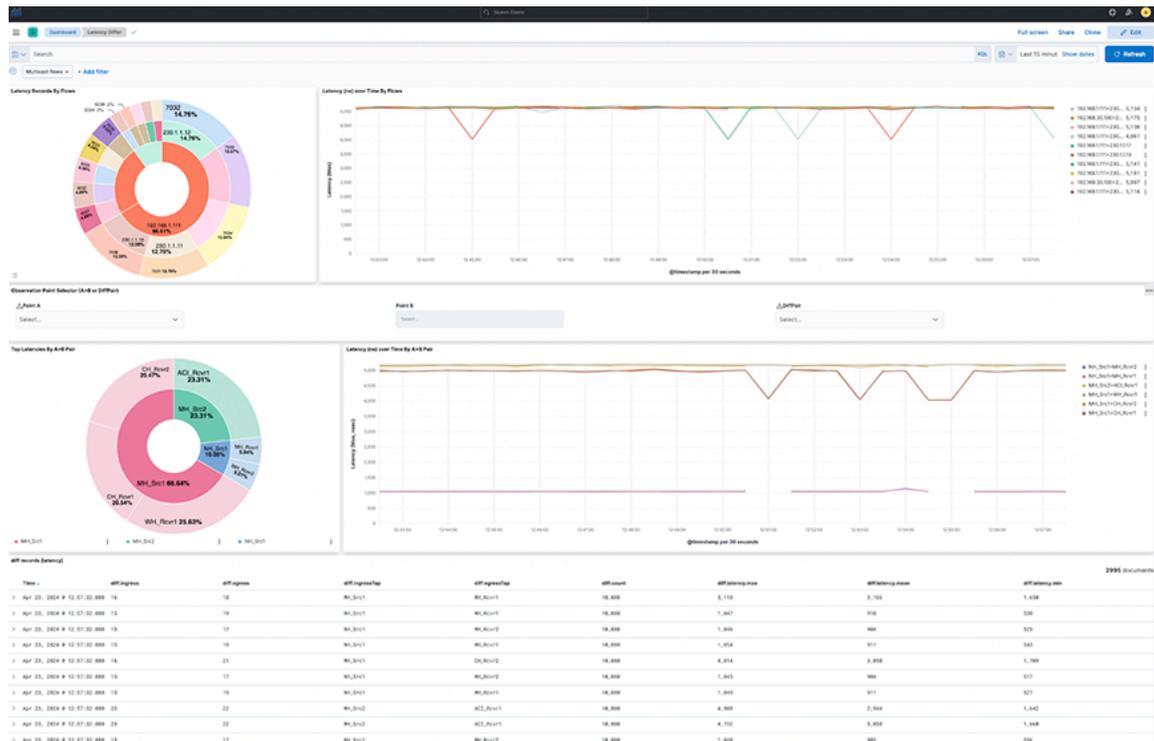
**Figure 5-1: Latency Differ Dashboard**



The following widgets appear in the Latency Differ dashboard:

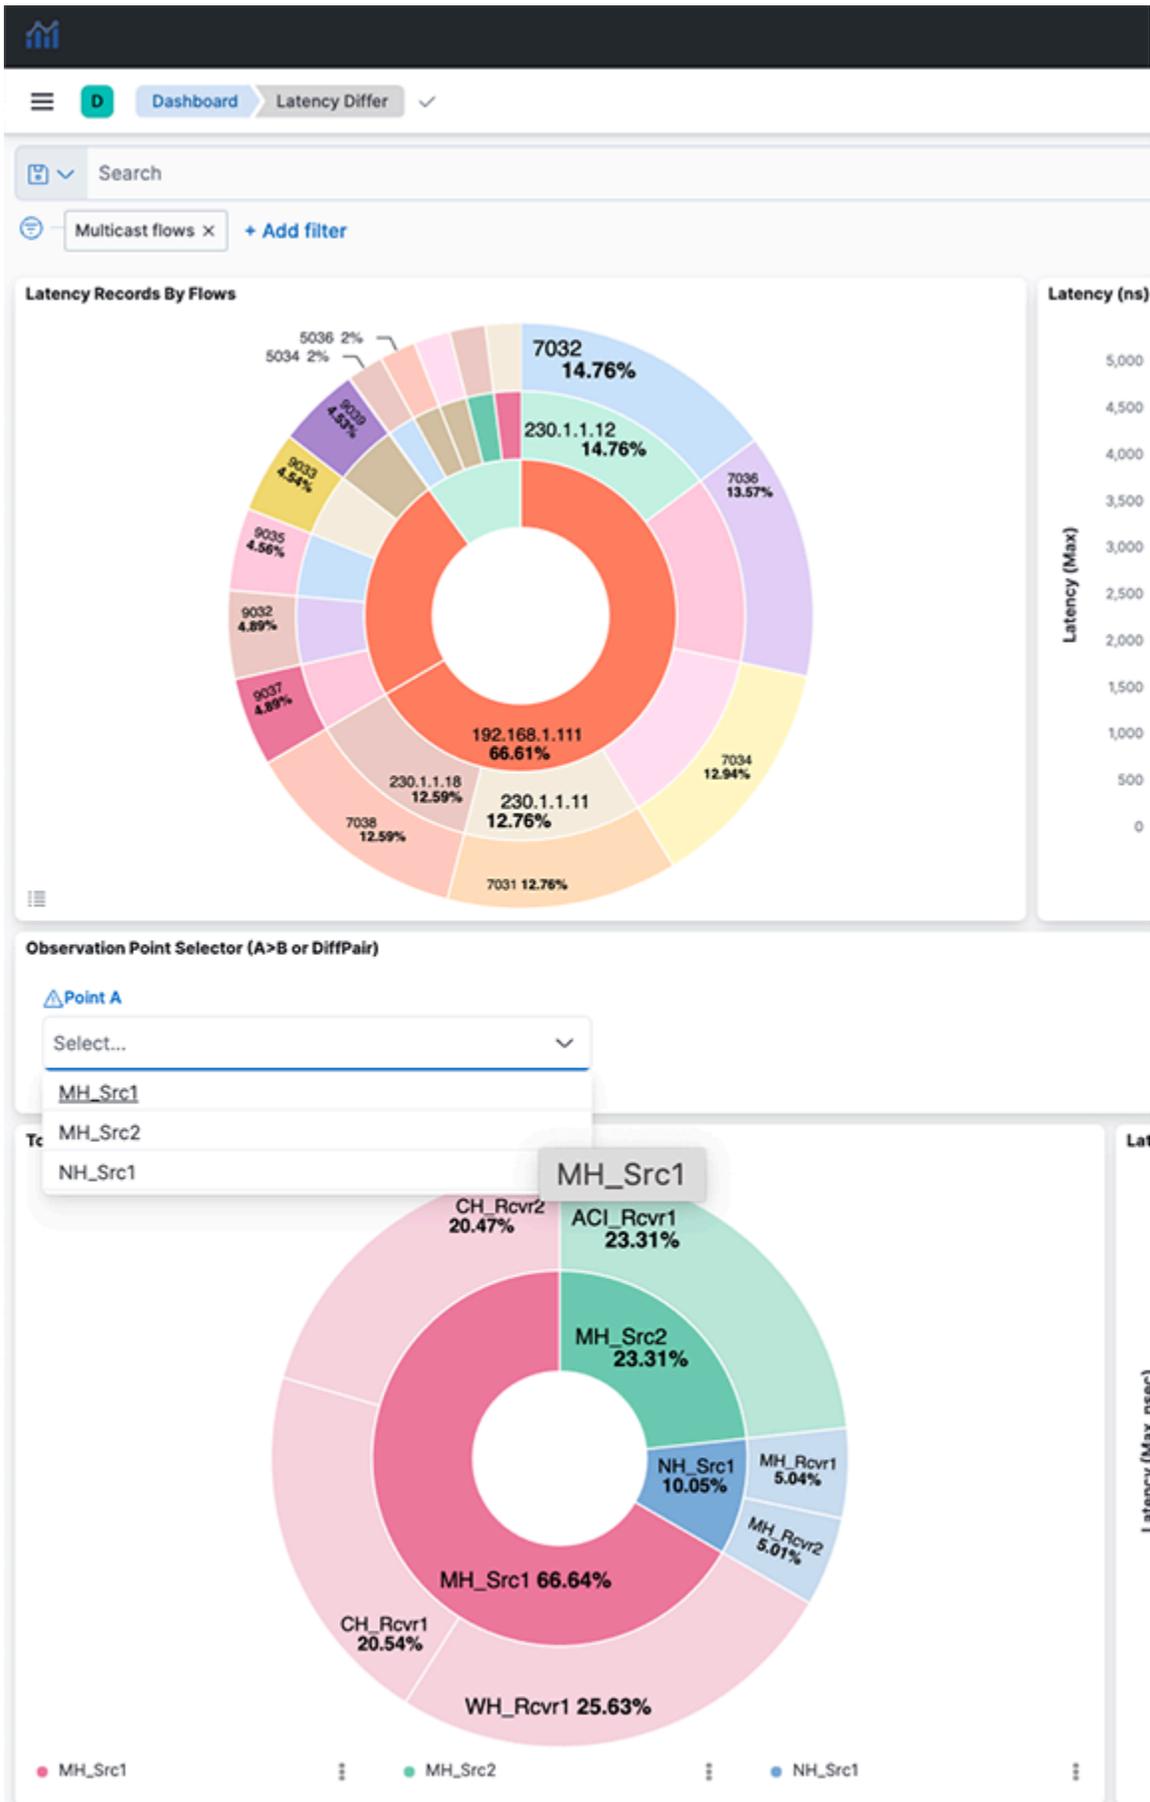
- **Latency Records By Flows:** The pie chart represents the proportions of flow latency summed. The inner circle displays source IP addresses, the middle circle displays destination IP addresses, and the outermost circle displays destination ports.
- **Latency over time By Flows:** The line chart represents the maximum Latency in nanoseconds (ns) over time, split by each flow between source IP and destination IP addresses.
- **Observation Point Selector (A > B or DiffPair):** Use the drop-down menus to filter by A > B pair or DiffPair. The point B selector is dependent on point A.
- **Top Latencies By A > B Pair:** The pie chart shows the Latency max summed by A > B Points. The inner circle displays the source A tap point, while the outer circle displays the B destination tap point.
- **Latency over time By A > B Pair:** The line chart represents maximum Latency in nanoseconds (ns) over time, split by each A > B pair between the source tap point and destination tap point.

Figure 5-2: Latency Record by Flows



Select **A > B** selection or **DiffPair** to visualize the data types. Filter the data using **A > B** Points by selecting a single source (**A**) and one or more receivers (**B**).

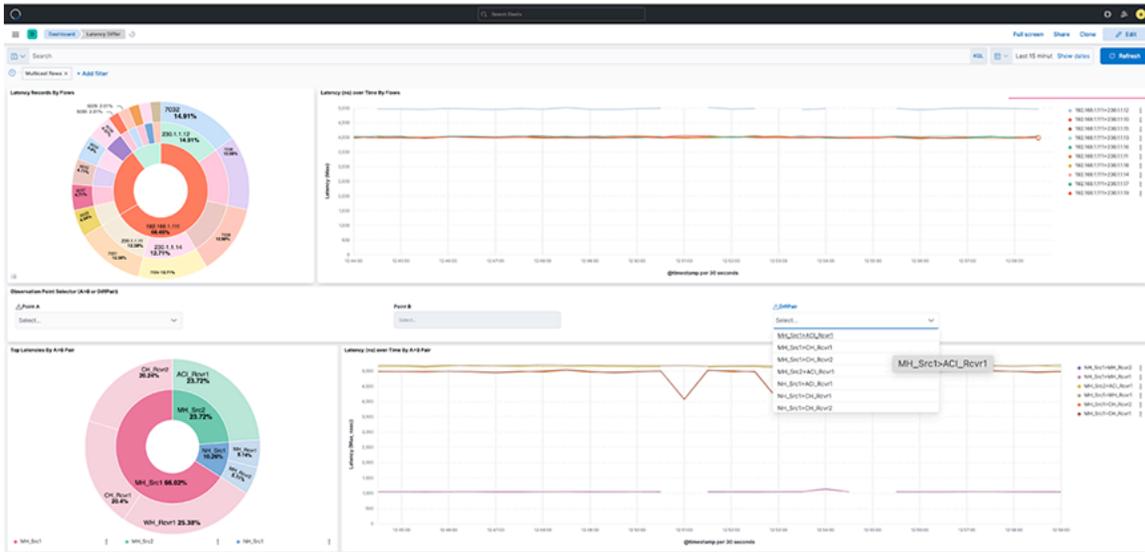
Figure 5-3: Flow Record with Observation Point Selector



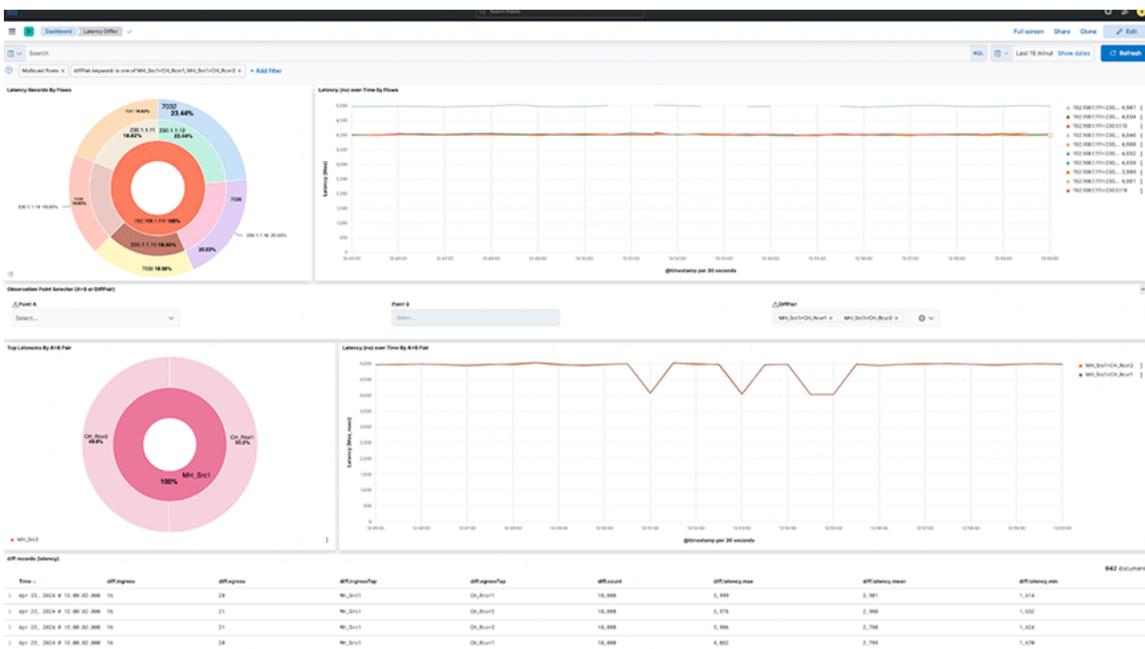


Change the visualization perspective by selecting DiffPairs by selecting one or more **DiffPair** for their analysis.

**Figure 5-6: DiffPair Analysis**



**Figure 5-7: Another DiffPair Analysis**

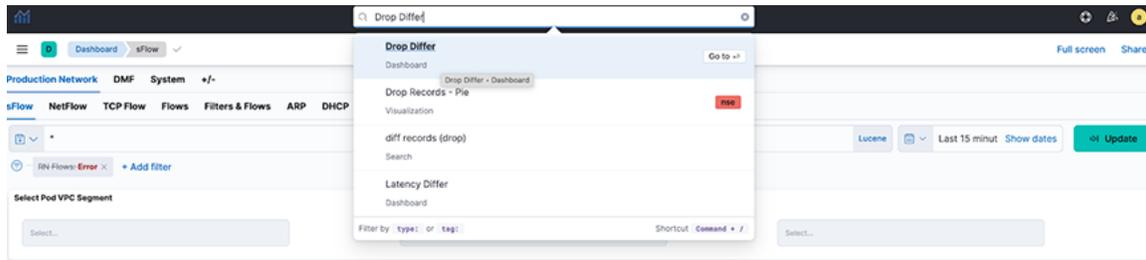


## Drop Differ Dashboard

Locate the **Drop Differ** dashboard by searching for the term **Drop Differ**.

The dashboard combines a visual representation of NetFlow Latency data in two views. The upper view displays individual flows, while the lower view aggregates A towards B pairs (**A > B**) or **DiffPair**. Drop Differ Dashboard

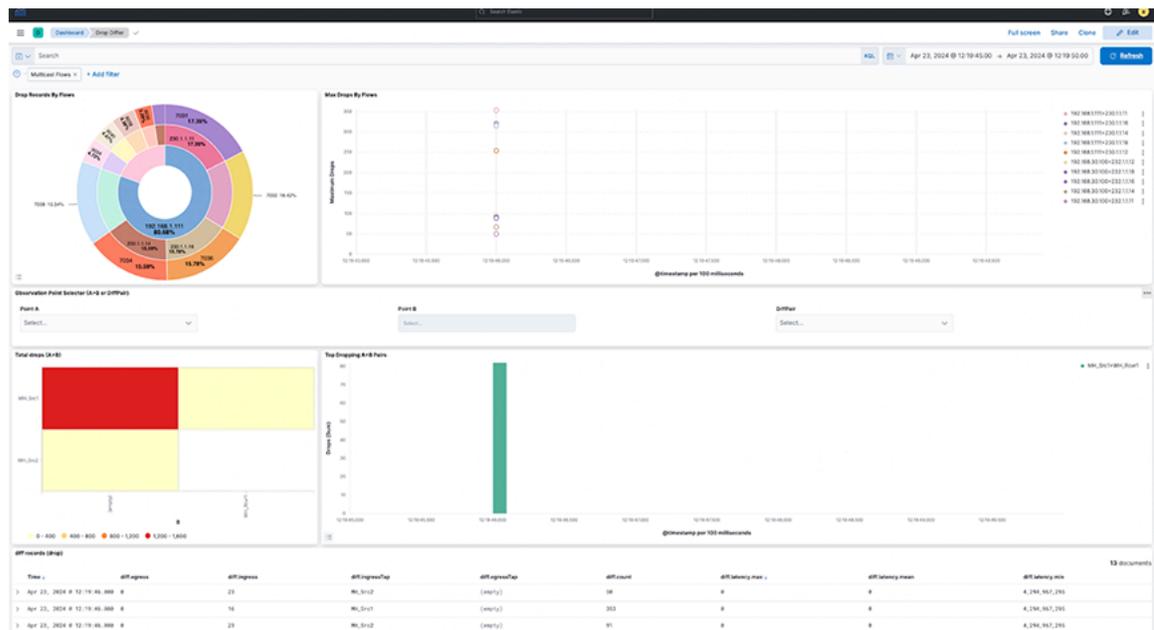
Figure 5-8: Drop Differ Dashboard



The following widgets appear in the Drop Differ dashboard:

- **Drop Records By Flows:** The pie chart represents the proportions of drop packets for each flow summed. The inner circle displays source IP addresses, the middle circle displays destination IP addresses, and the outermost circle displays destination ports.
- **Max Drops By Flows:** The line chart represents the maximum number of drop packets, separated by each flow between source IP and destination IP addresses. If fewer data points exist, the chart displays them as individual points instead of complete lines.
- **Observation Point Selector (A>B or DiffPair):** Use the drop-down menus to filter by **A > B** pair or **DiffPair**. The point **B** selector is dependent on point **A**.
- **Top Drop A>B:** The heat map displays the drop of packets summed by **A > B** Points. The map plots the source tap point, **A** on the vertical axis and the destination tap point, **B**, on the horizontal axis.
- **Top Dropping A>B Pairs:** The bar chart represents the sum of drop packets over time, separated by each **A > B** pair between the source tap point and the destination tap point. It shows the **Top 10** available dropping **A > B** pairs .

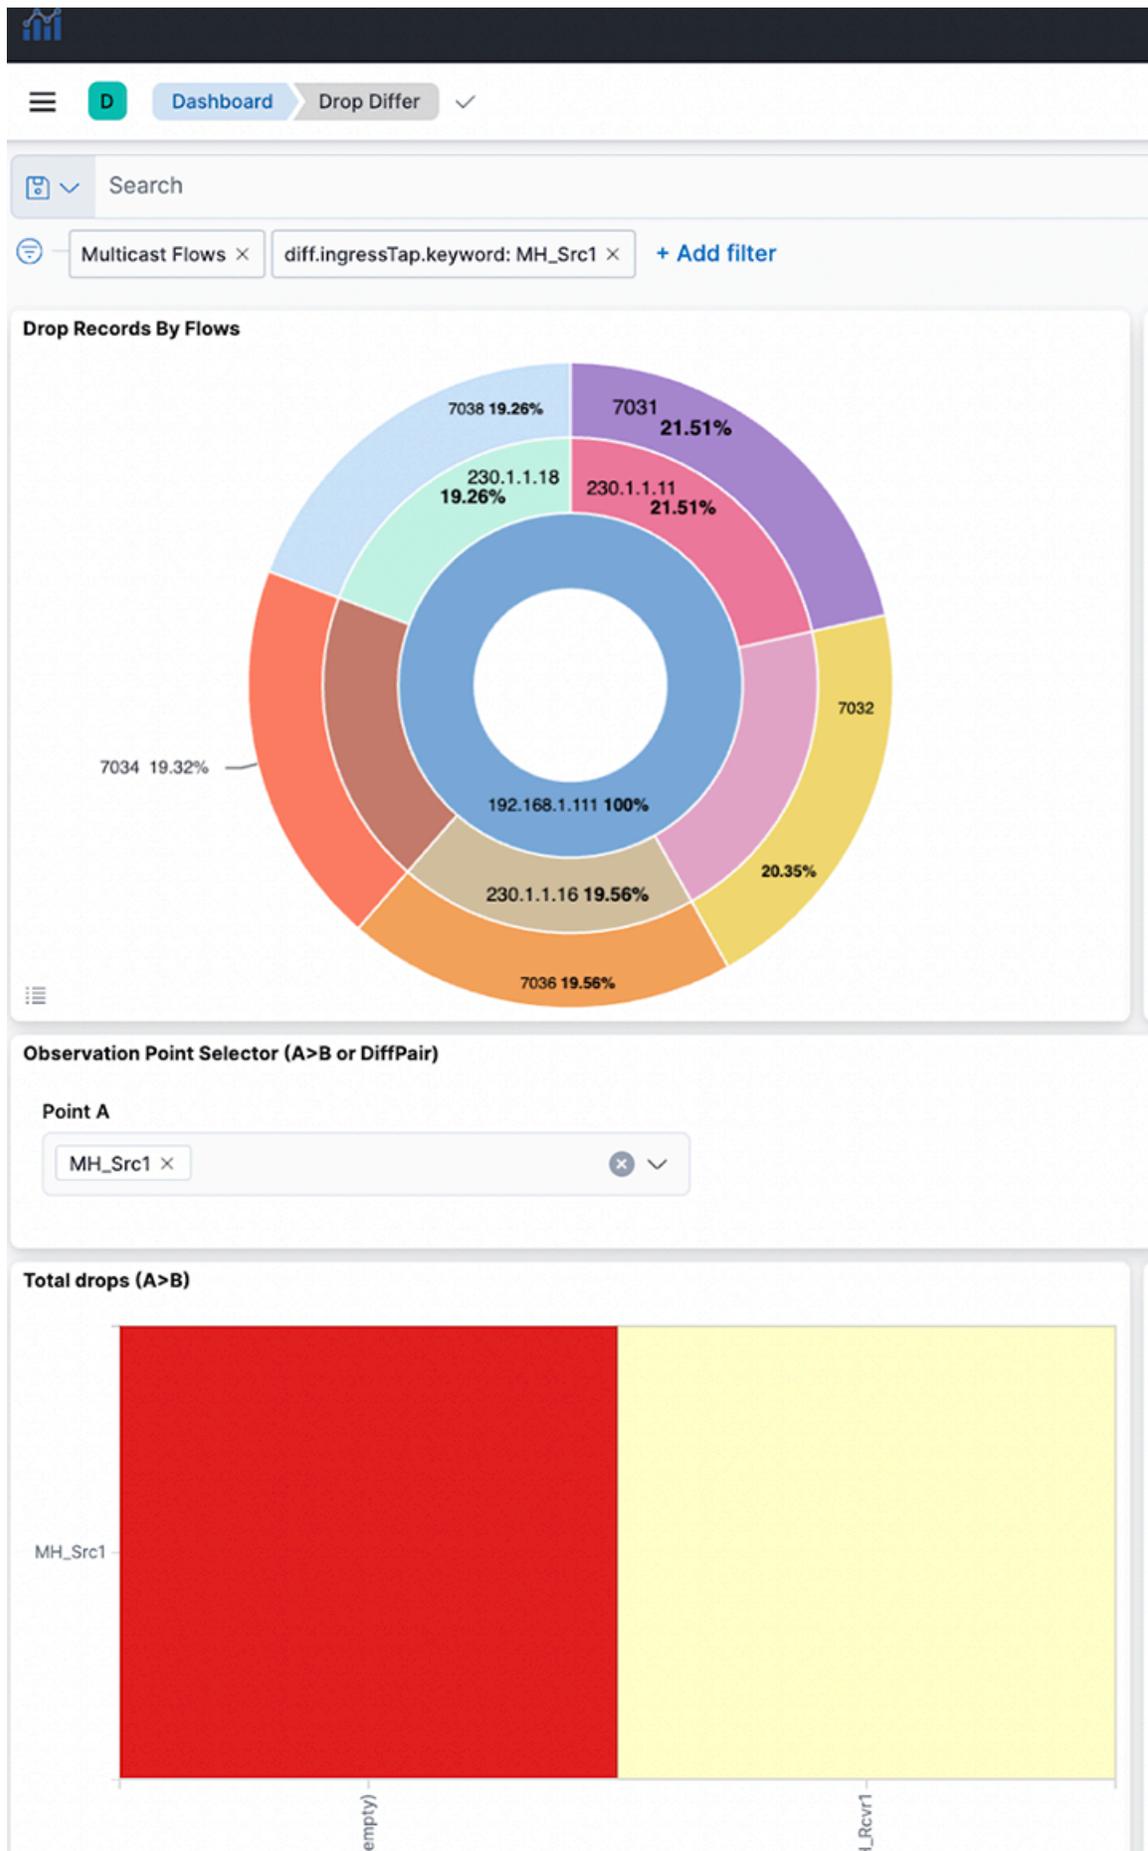
Figure 5-9: Top Dropping A>B Pairs



Select **A > B** selection or **DiffPair** to visualize the data types.

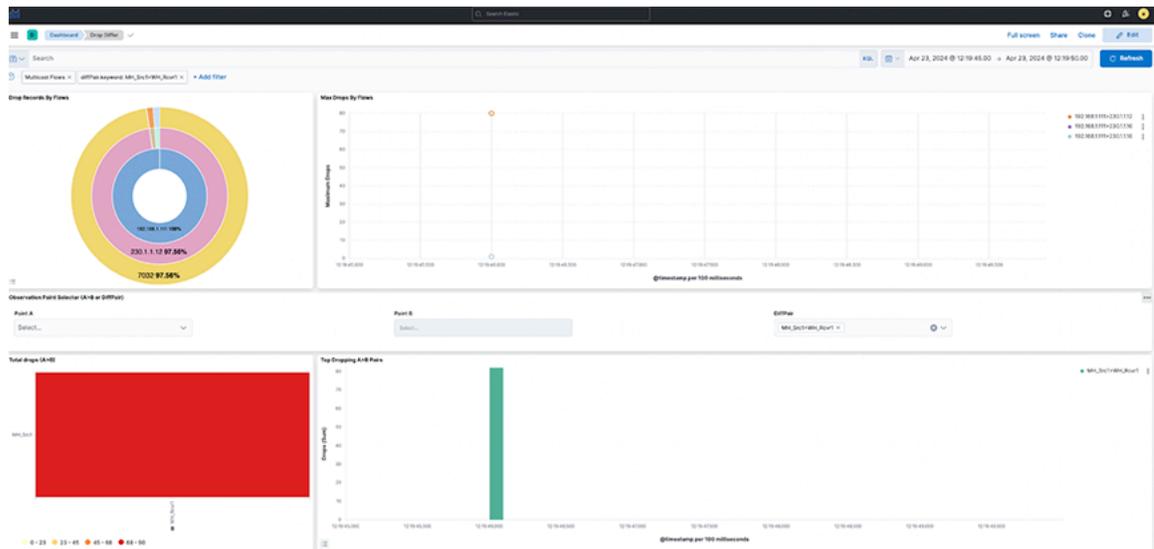
Filter the data using **A > B** Points by selecting a single source (**A**) and one or more receivers (**B**).

**Figure 5-10: Data Types Visualization**



- This provides a dashboard for packet drops between points A and B(s), either split by flows in between those points (Top) or filtered by **A > B** pairs (bottom) as selected. View the diff records at the bottom of the dashboard.
- Select individual data points in the visualization for further analysis.
- Selecting DiffPairs can provide a similar visualization perspective. Choose one or more DiffPairs for analysis.

**Figure 5-12: DiffPair Analysis for Drop Differ**



### Configuring Watcher Alerts

Watcher is an elastic search feature that supports the creation of alerts based on conditions triggered at set intervals. For more information, refer to: [Watcher | Kibana Guide \[7.17\] | Elastic](#)

AN includes two built-in examples of watcher templates for ease of use. To access the templates, navigate to **Stack Management > Watcher**.

- Arista\_NetOps\_Drop\_Differ\_Watch
- arista\_NetOps\_Latency\_Differ\_Watch

The templates are disabled by default and require manual configuration before use.

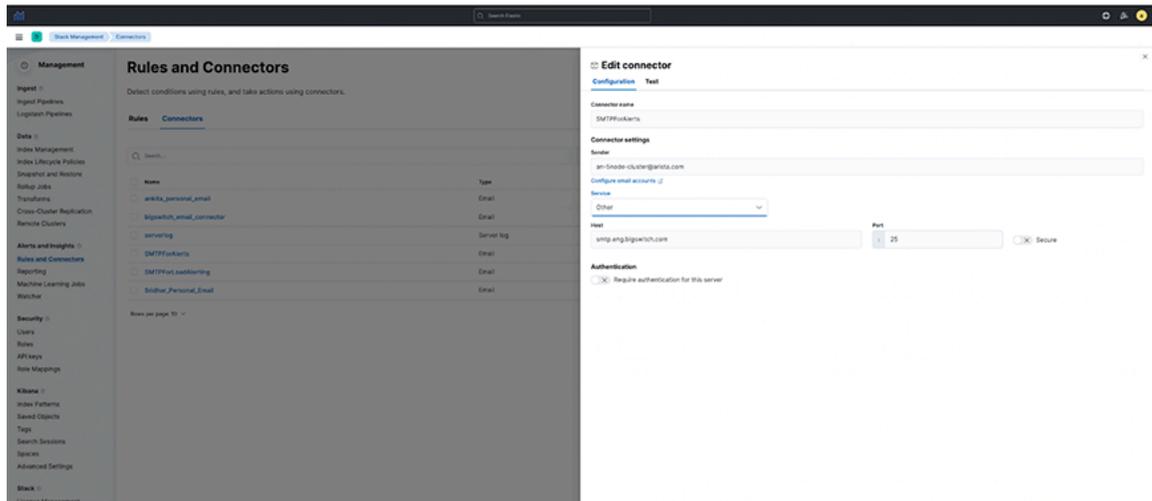
### Setting the SMTP Connector

The system dispatches **Alerts** by email; configure the **SMTPForAlerts Connector** before use.

1. Navigate to **Stack Management > Connector**.
2. Under **Configuration** for the **SMTPForAlerts Connector**, specify the **Sender** and **Service** field values.
3. Sending email alerts may require authentication based on the type of mail service selected.

4. Test and validate the settings using the **Test** tab.

**Figure 5-13: Testing SMTP Connector**



### Setting the Watchers

- **arista\_NetOps\_Drop\_Differ\_Watch:**

1. The watcher is configured to send an alert when the maximum drop count of packets in **NetFlow** in the last 5-minute interval exceeds the historical average (last 7-day average) of drop of packets by a **threshold percentage**.
2. This watcher is configured by default to be **triggered** every 10 minutes.
3. As this may be incorrect for all flows combined, configure it for a particular **Flow** and **Destination Port**.
4. Search for **CHANGE\_ME** in the watcher and specify the flow and destination port value (introduced to correctly compare each flow and destination port individually instead of comparing all flows together).
5. Specify the percentage\_increase parameter in the condition using a positive value between **0-100**.
6. Enter the recipient's email address receiving the alert.

7. Select **Save watch**.

Figure 5-14: NetOps\_Drop\_Differ\_Watch-1

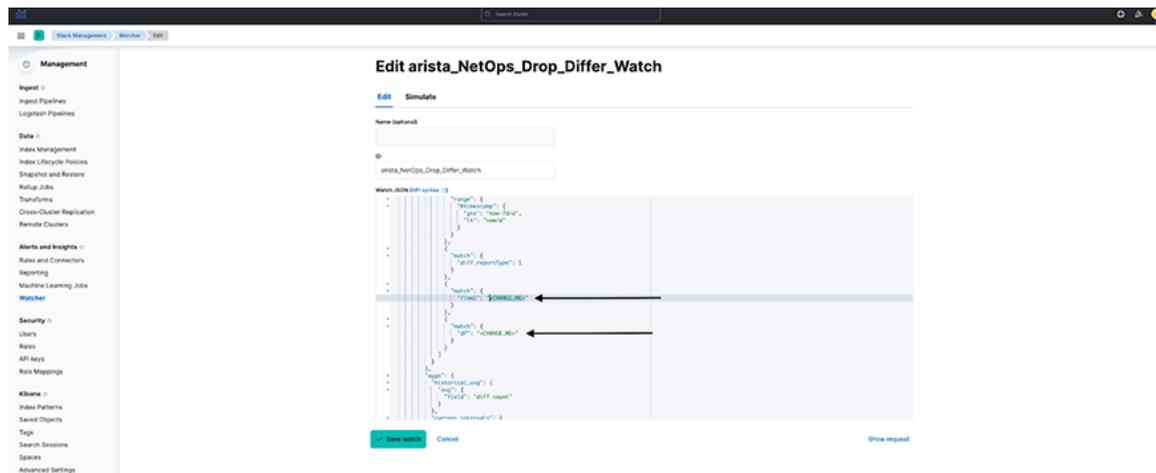


Figure 5-15: NetOps\_Drop\_Differ\_Watch-2

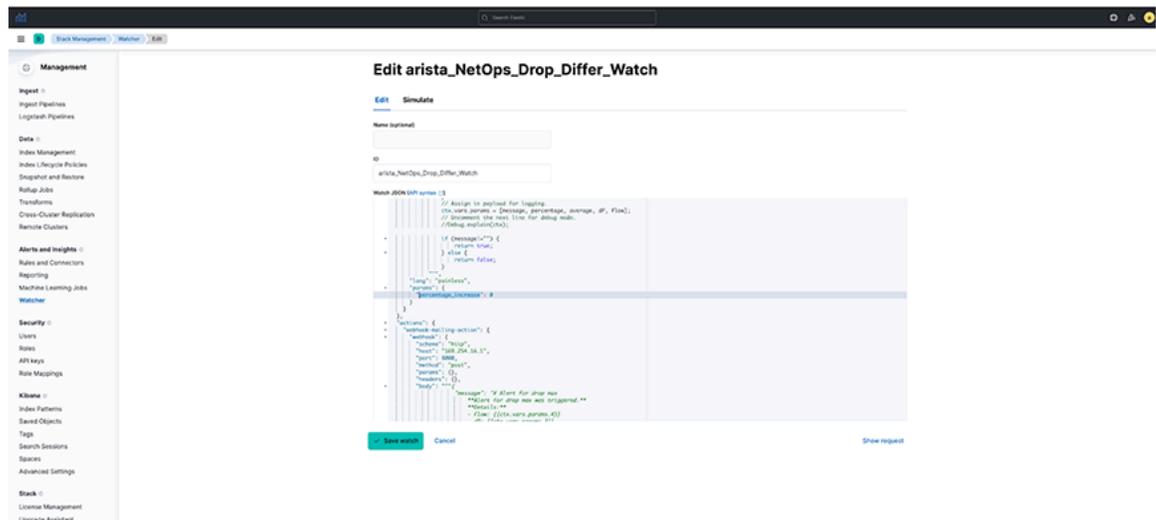
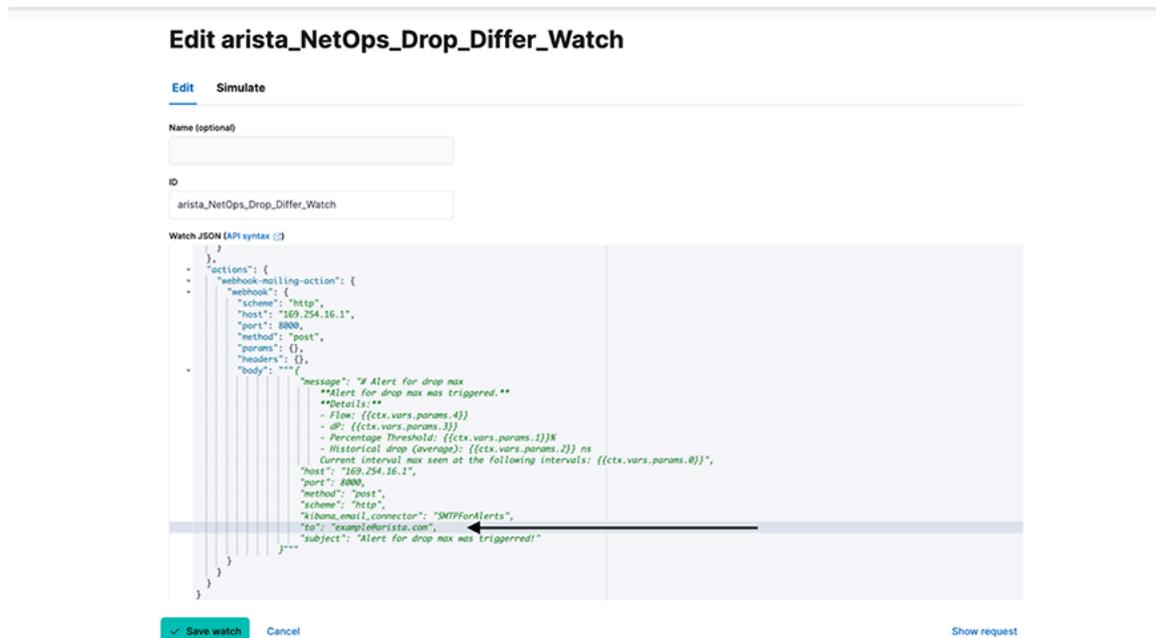


Figure 5-16: Editing NetOps\_Drop\_Differ\_Watch



- **arista\_NetOps\_Latency\_Differ\_Watch:**

1. The watcher is configured to send an alert when **NetFlow**'s maximum latency (or lag) in the last 5-minute interval exceeds the historical average (last 7-day average) latency by a **threshold percentage**.
2. This watcher is configured by default to be **triggered** every 10 minutes.
3. As this may be incorrect for all flows combined, configure it for a particular **Flow** and **Destination Port**.
4. Search for **CHANGE\_ME** in the watcher and specify the flow and destination port value (introduced to correctly compare each flow and destination port individually instead of comparing all flows together).
5. Specify the `percentage_increase` parameter in the condition using a positive value between **0-100**.
6. Enter the recipient's email address receiving the alert.
7. Select **Save watch**.

## Considerations

- Default Watchers are disabled and must be modified with user-configured alert settings before being enabled.

**Figure 5-17: Arista\_NetOps\_Drop\_Differ\_Watch**

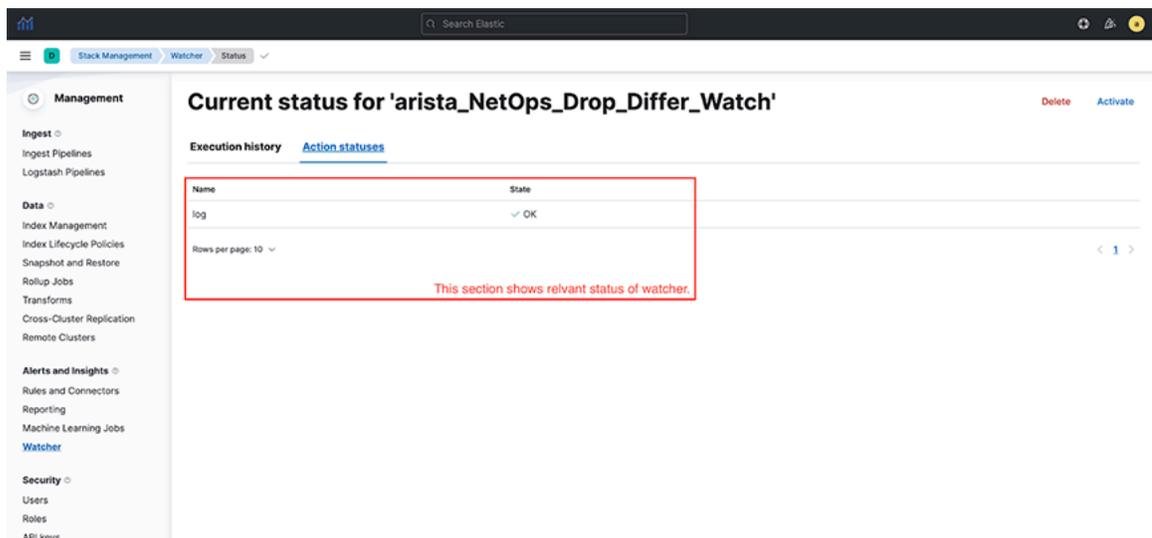


## Troubleshooting

- The dashboard obtains its data from the **flow-netflow** index. If no data is present in the dashboard, verify there is sufficient relevant data in the index.

- Watchers trigger at a set interval. To troubleshoot issues related to watchers, navigate to **Stack Management > Watcher**. Select the requisite watcher and navigate to **Action statuses** to determine if there is an issue with the last trigger.

**Figure 5-18: Watcher Action Status**



### Usage Notes

- The dashboards only show partial and not full drops during a given time and are configured with filtering set to the **egress.Tap** value as **empty**.
- A **full drop** occurs when the flow of packets is observed at the source tap point, but no packet is observed at the destination tap point. The dashboards are configured to filter out full drop flows.
- A **partial drop** is a scenario in which the flow of packets is observed at the source tap point, and some, if not all, packets are observed at the destination tap point. The dashboards clearly show partial drop flows.

## Monitoring DMF Network Health

---

This chapter describes uses the dashboards on the **DMF Network** tab to monitor activity on the DANZ Monitoring Fabric. It includes the following sections.

- [DMF Network Tab](#)
- [Policy Statistics Dashboard](#)
- [Interface Statistics](#)
- [SN \(Service Node\) Statistics](#)
- [Events](#)

### 6.1 DMF Network Tab

The **DMF Network** tab includes dashboards that display the following information visible to the DMF controller:

- Policy Statistics
- Interface Statistics
- SN Statistics
- Incline Statistics
- Events

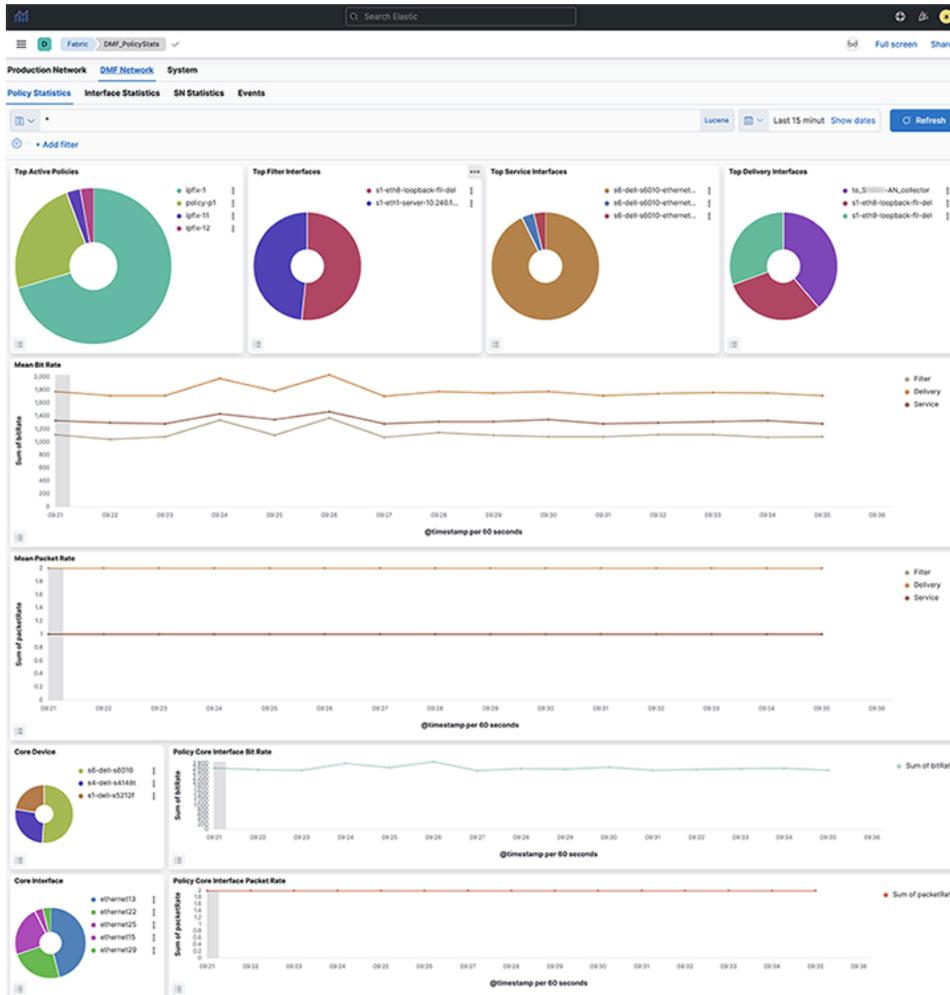


**Note:** Information displayed on these dashboards requires configuring an ACL for Redis and replicated Redis using the Analytics CLI after first boot configuration.

## 6.2 Policy Statistics Dashboard

Click the **Policy Statistics** tab to display the following dashboard:

**Figure 6-1: DMF Network > Policy Statistics Dashboard**



The **Policy Statistics** dashboard summarizes information about DANZ Monitoring Fabric policy activity and provides the following panels:

- Top Active Policies
- Top Filter Interfaces
- Top Service Interfaces
- Top Delivery Interfaces
- Mean Bit Rate
- Mean Packet Rate
- Core Switch
- Policy Core Interface Bit Rate
- Core Interface
- Policy Core Interface Packet Rate
- Records
- Policies with no traffic

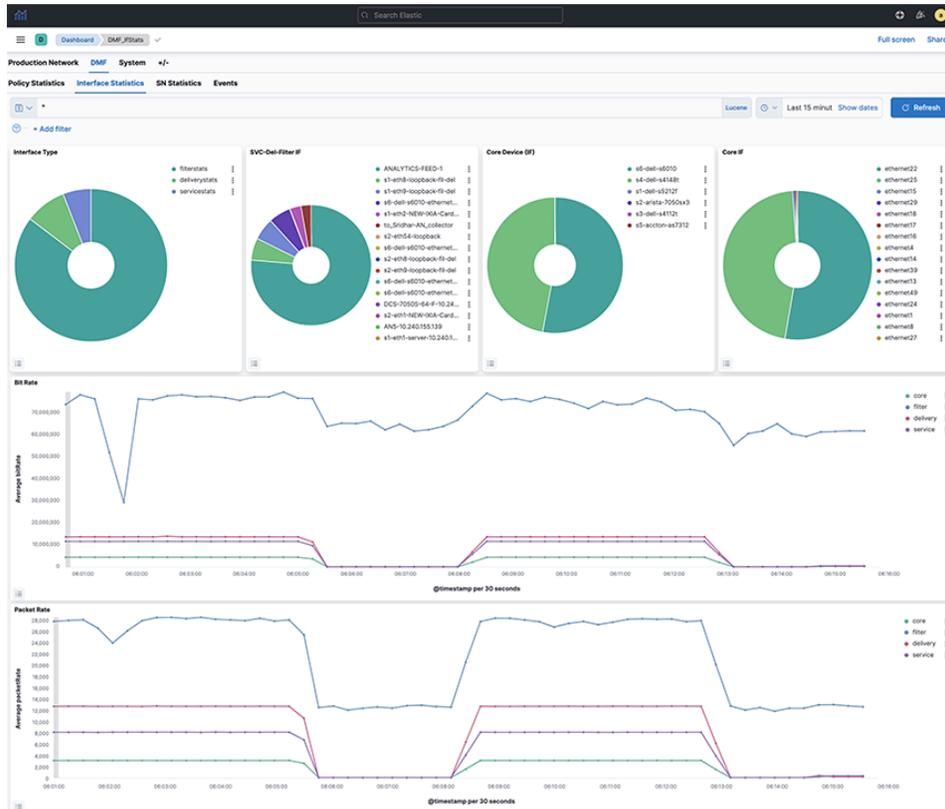
Use the **Top Active Policies** visualization to verify that your DANZ Monitoring Fabric policies are active and behaving as expected.

Use the **Filter Interfaces** visualization to balance the utilization of your filter interfaces and ensure that it doesn't drop any packets to analyze.

## 6.3 Interface Statistics

Click the **Interface Statistics** tab to display the following dashboard.

Figure 6-2: DMF Network > Interface Statistics



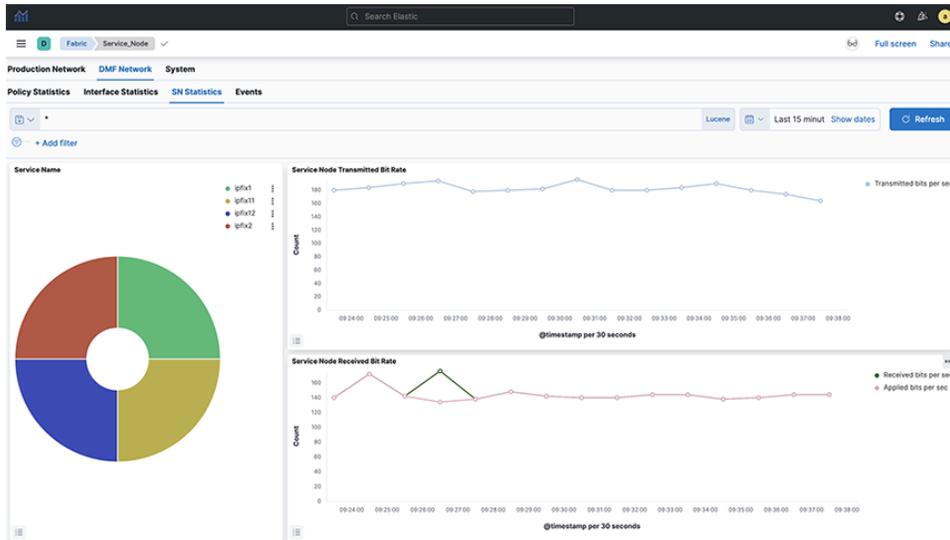
The **Interface Statistics** dashboard summarizes information about DANZ Monitoring Fabric switch interface activity and provides the following panels:

- Interface Type
- SVC-Del-Filter IF
- Core Switch (IF)
- Core IF
- Bit Rate
- Packet Rate
- Interface Detail

## 6.4 SN (Service Node) Statistics

Click the **SN Statistics** tab to display the following dashboard.

Figure 6-3: DMF Network > SN Statistics

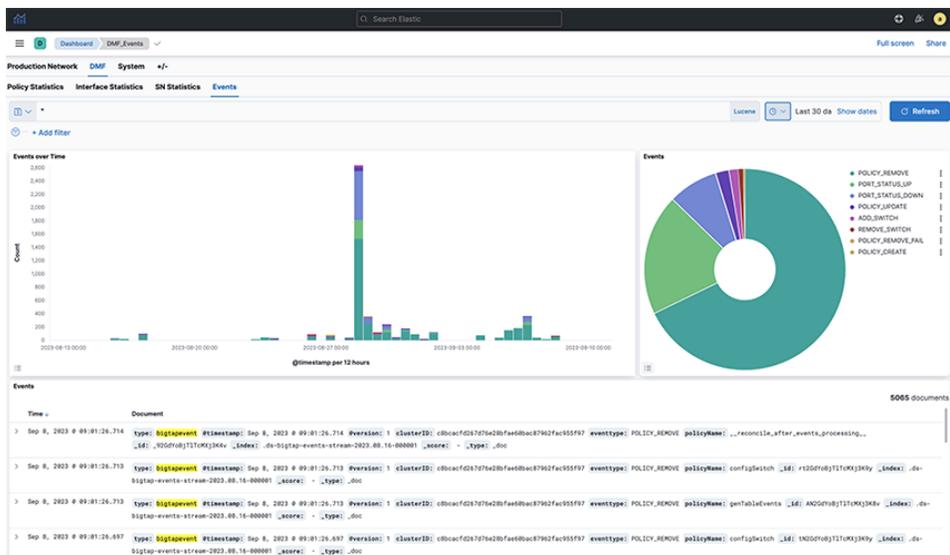


Select a service from the pie chart to see the statistics for a specific managed service. It will display statistics for the selected service.

## 6.5 Events

Click the **Events** tab to display the following dashboard.

Figure 6-4: DMF Network > Events



The **Events** dashboard summarizes information about DANZ Monitoring Fabric management network events and provides the following panels:

- Events Over Time

- Events

## Monitoring Users and Software Running on the Network

This chapter describes using Arista Analytics with the DMF Recorder Node. It includes the following sections.

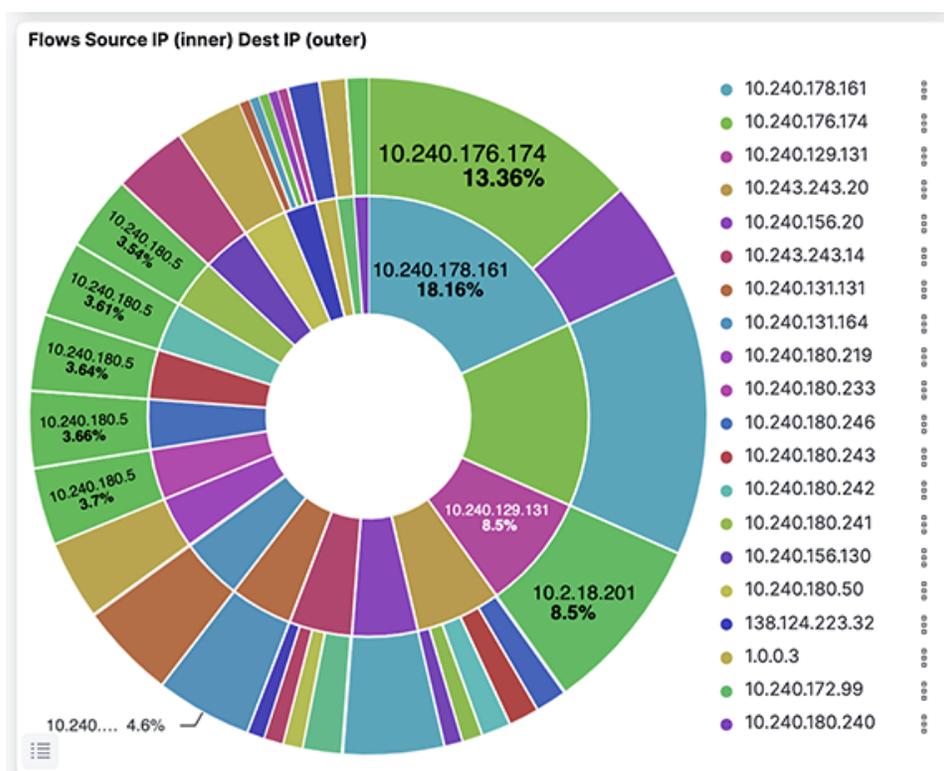
- [IP Addresses](#)
- [Geographic Location](#)
- [Software Running in the Network](#)
- [User Activity](#)
- [Monitoring Active Directory Users](#)

### 7.1 IP Addresses

This section describes identifying traffic transmitted or received by the source or destination IP address.

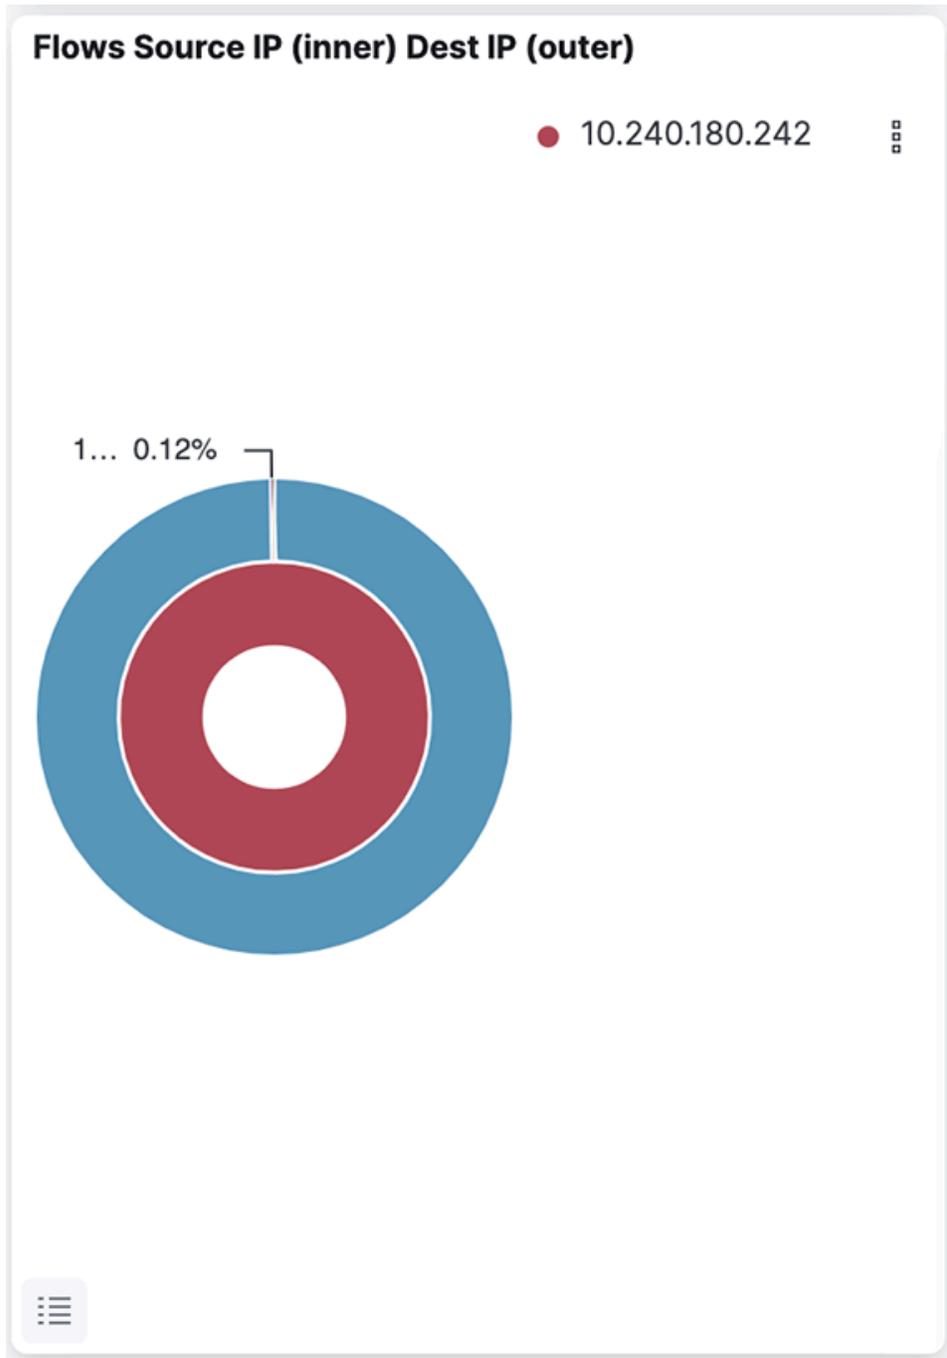
#### 7.1.1 Source and Destination Addresses

Figure 7-1: Identifying Source and Destination IP Addresses



Click an IP address, then click the **Magnifying Glass** icon (+) to pin the address to the dashboard.

**Figure 7-2: Filtering Results by IP Address**



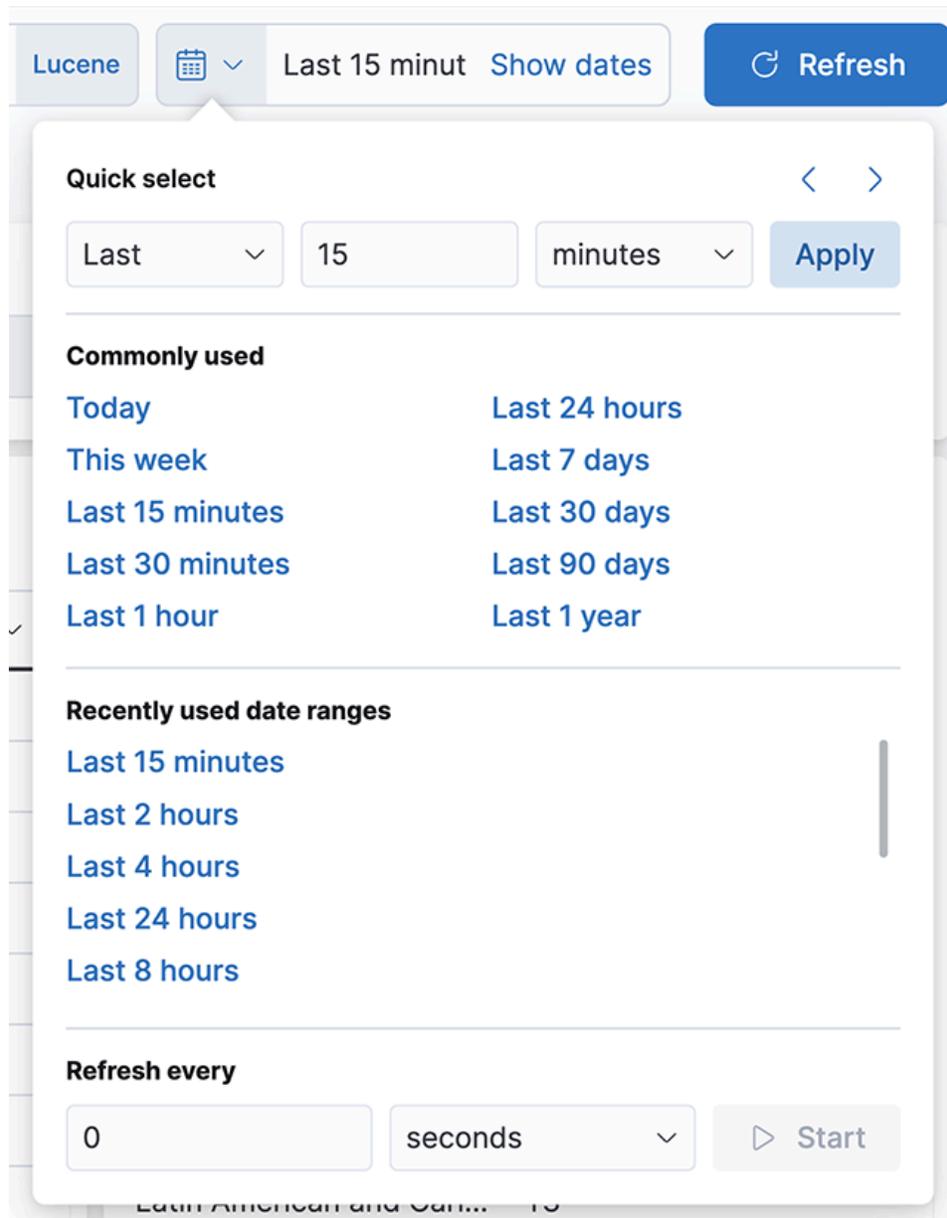
The selected IP address is added to the filters on the dashboard.

Each dashboard has a bar chart depicting traffic on the y-axis and time on the x-axis. To add a time filter, click and drag an area in the **All Flows Over Time** bar chart.

## 7.1.2 Unauthorized IP Destinations

To determine if an IP destination that is not authorized is being accessed in your network for a specific period, set the time value in the upper right corner.

**Figure 7-3: Setting the Duration**



Select the duration of time for the search.

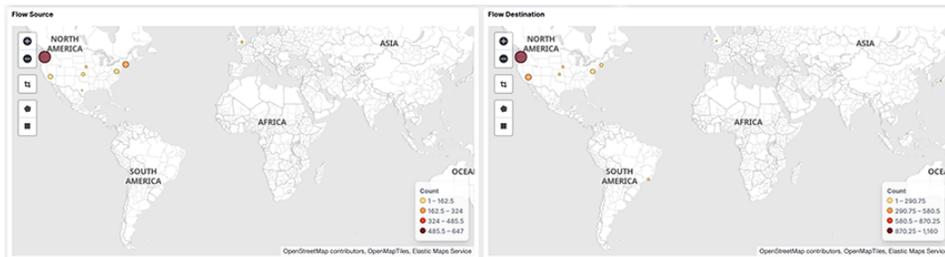
Type the IP address or the Network ID in the **Search** field.

The system displays any events associated with the address or network ID.

## 7.2 Geographic Location

Analytics associates public network IP addresses to geographic regions using the MaxMind GeoIP database. Traffic associated with these addresses shows as a heat map on the **Map** visualization on the **sFlow**<sup>®</sup> dashboard. To filter on a region, draw a box or a polygon around the region.

**Figure 7-4: Geographic Flow Source and Destination**



Use the **Square** tool to draw a square around a region of interest, or use the **Polygon** tool to draw an irregular shape around a region. It will redraw the map to zoom in on the selected region and to show details about traffic to or from the region.

## 7.3 Software Running in the Network

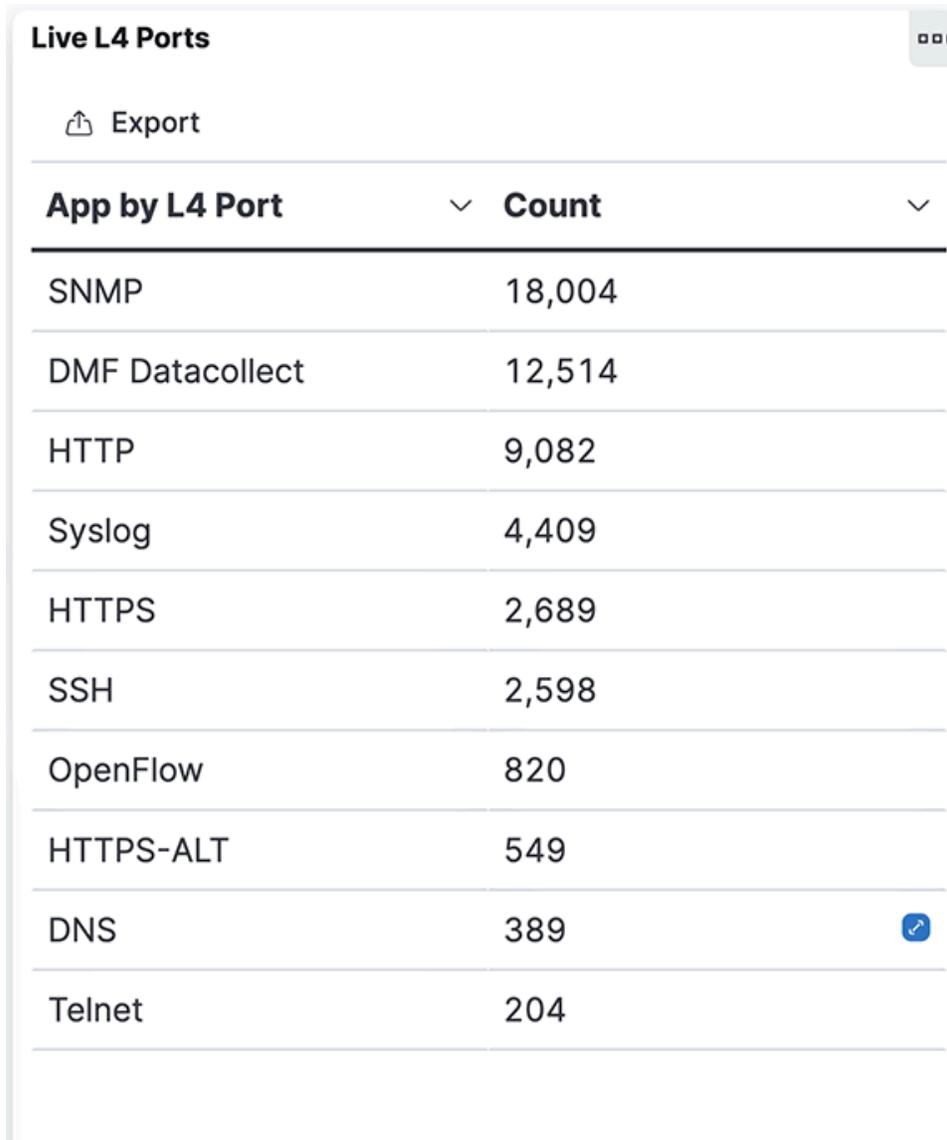
This section identifies specific applications or operating systems running on network hosts.

\* sFlow<sup>®</sup> is a registered trademark of Inmon Corp.

### 7.3.1 Top Talkers Using Well-known Layer-4 Ports

To view top-N statistics for the flows using a well-known L4 port, use the **Live L4 Ports** table on the **Flows** dashboard.

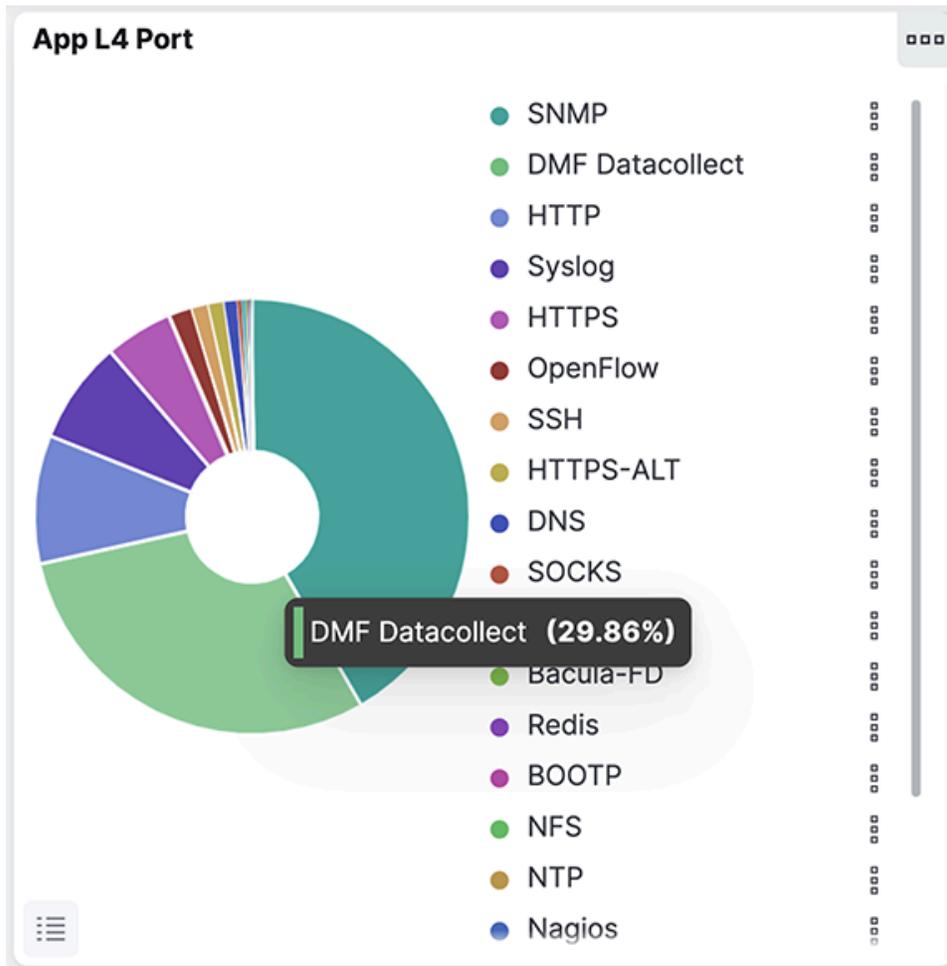
Figure 7-5: Flows > Live L4 Ports



App by L4 Port	Count
SNMP	18,004
DMF Datacollect	12,514
HTTP	9,082
Syslog	4,409
HTTPS	2,689
SSH	2,598
OpenFlow	820
HTTPS-ALT	549
DNS	389
Telnet	204

Use the **App L4 Port** table on the **sFlows** dashboard when a sFlow generator configured to send flows to Analytics.

**Figure 7-6: sFlow > App L4 Port**



These tables use well-known ports to identify the traffic generated by each application. You can also associate user-defined ports with applications as described in the following section.

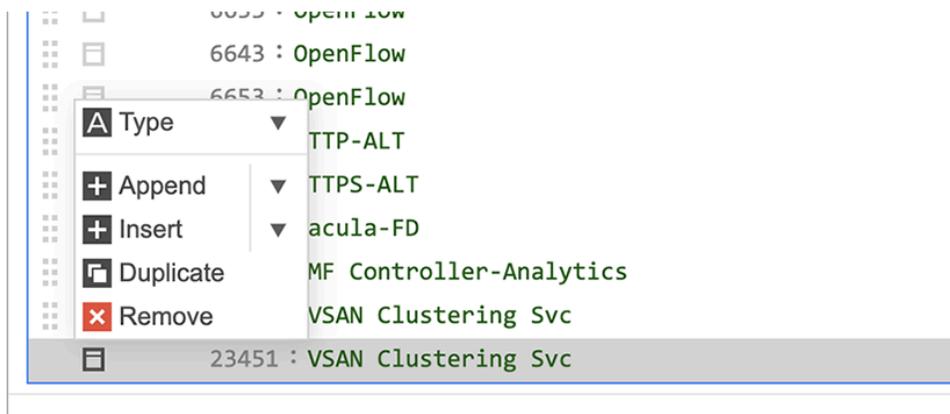
### 7.3.2 Associating Applications with User-defined Layer4 Ports

To associate user-defined ports with applications, complete the following steps:

1. Select **System > Configuration**.

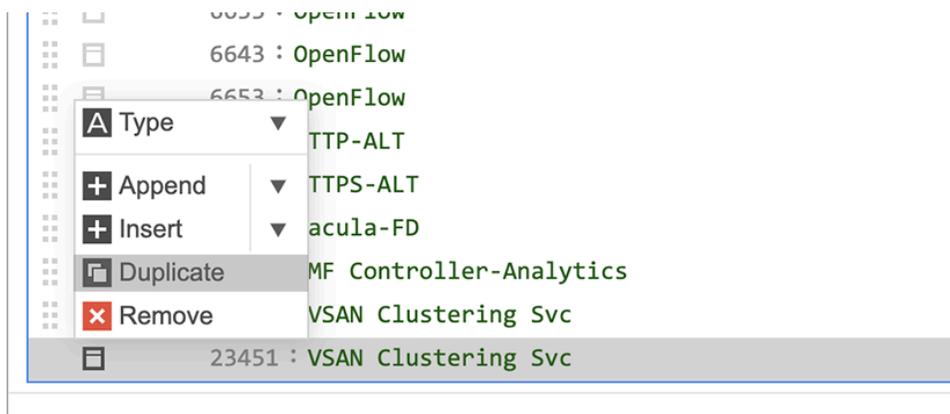
2. Select the **Edit** control to the right of the **Ports** section.

**Figure 7-7: Edit Ports**



3. To copy an existing row, enable the checkbox to the left of the row and select **Duplicate** from the drop-down menu.

**Figure 7-8: Duplicate Ports**



4. Type over the port number in the row you copied and enter an associated label.  
For example, assign **port 1212** to **Customer App X**.
5. Click **save**.

### 7.3.3 Software Running on Hosts

The following features identify the software running on hosts in the monitored network.

- Searching for well-known applications
- Using Layer4 labels
- Searching packet captures on the DMF Recorder Node
- Using the **Flows** dashboard
- Using the **DHCP** dashboard for information about operating systems

The IP block default mapping associates many common applications with specific address ranges. For example, you can identify video traffic by searching for **YouTube** or **Net#ix**.

L4 label strings identify applications using well-known ports and applications running on user-defined ports after mapping those ports to the applications.

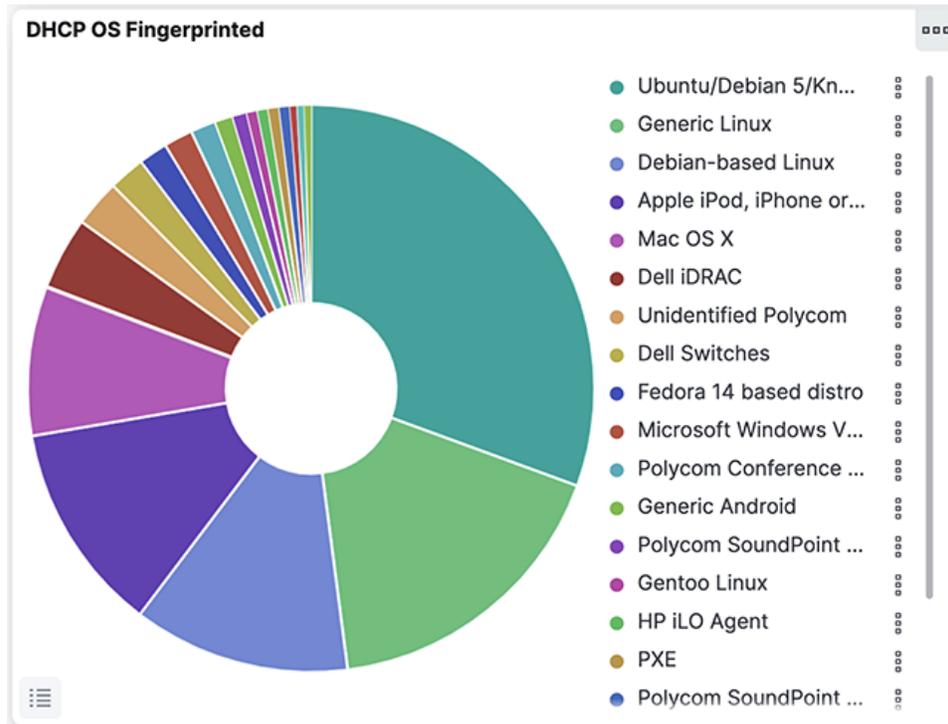
The flow dashboards all give an overall sense of who is talking to whom. Click on an IP address or L4 port, and with the **+** that appears, pin that to filter the dashboard by the selection. Every dashboard has a bar

chart depicting traffic on the y-axis and time on the x-axis. Note that a time filter can be added by a click and sideways selection of the bar chart.

The who can also be in terms of the user with a source of users to IP mappings (OpenVPN supported) configured. After that, a search by the user string for traffic attributed to that user over a dashboard period.

The DHCP dashboard indicates the operating systems running on hosts based on information derived from DHCP client requests. The default mapping is copied from the signatures provided by [fingerbank.org](http://fingerbank.org).

**Figure 7-9: DHCP OS Fingerprinting**



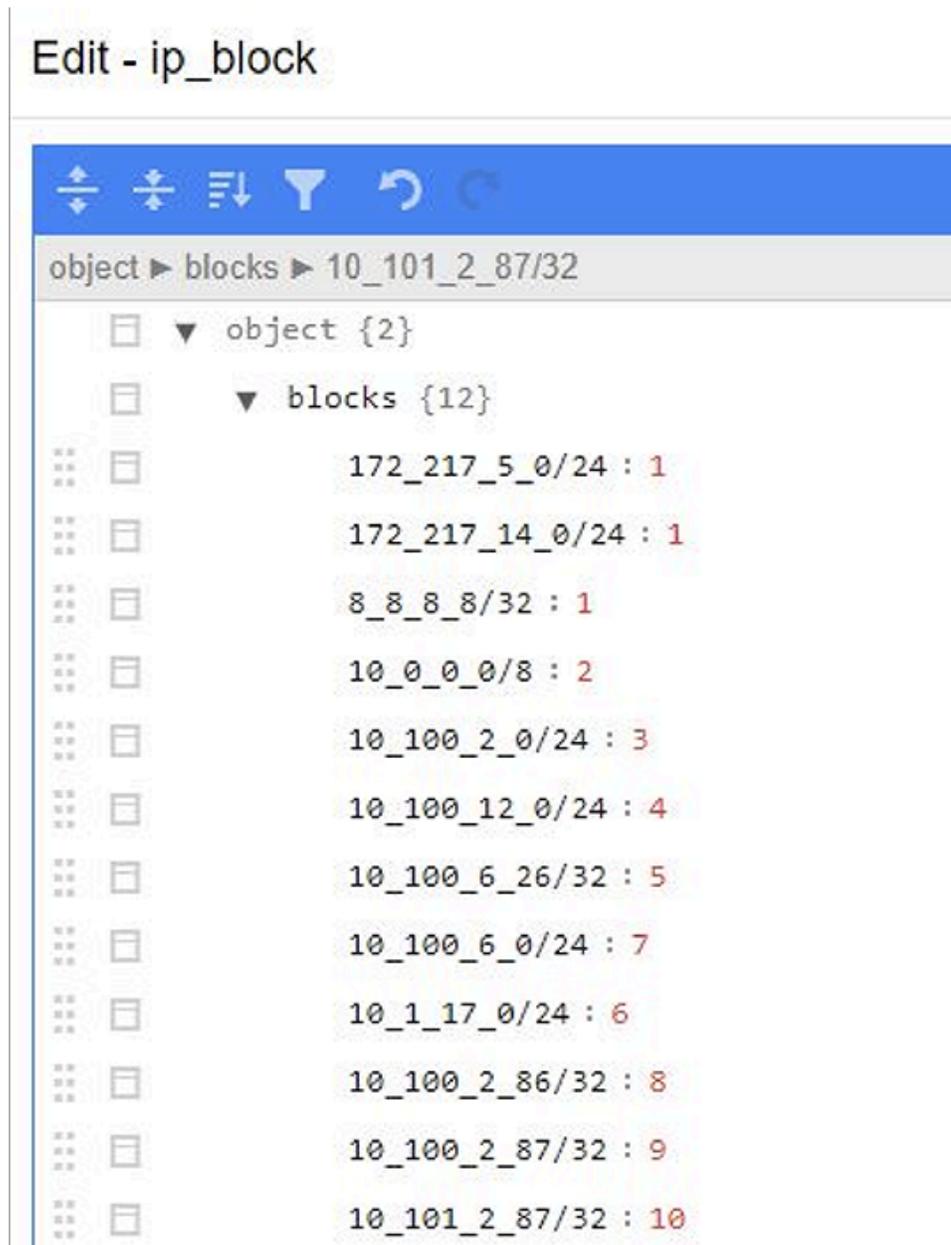
### 7.3.4 Tools Receiving Traffic

Identify traffic forwarded to a specific tool or host using the IP Blocks mapping to associate an IP address or a range of IP addresses to a label describing the application. This label will then appear on any dashboards or visualizations that display the IP Block labels. After mapping, the search can happen for events associated with the label assigned to the tool.

Refer to the [Mapping IP Address Blocks](#) section for details about updating the IP block mapping file.

1. To edit the IP blocks, select **System > Configuration** and click the **Edit** control to the right of the IP blocks section.

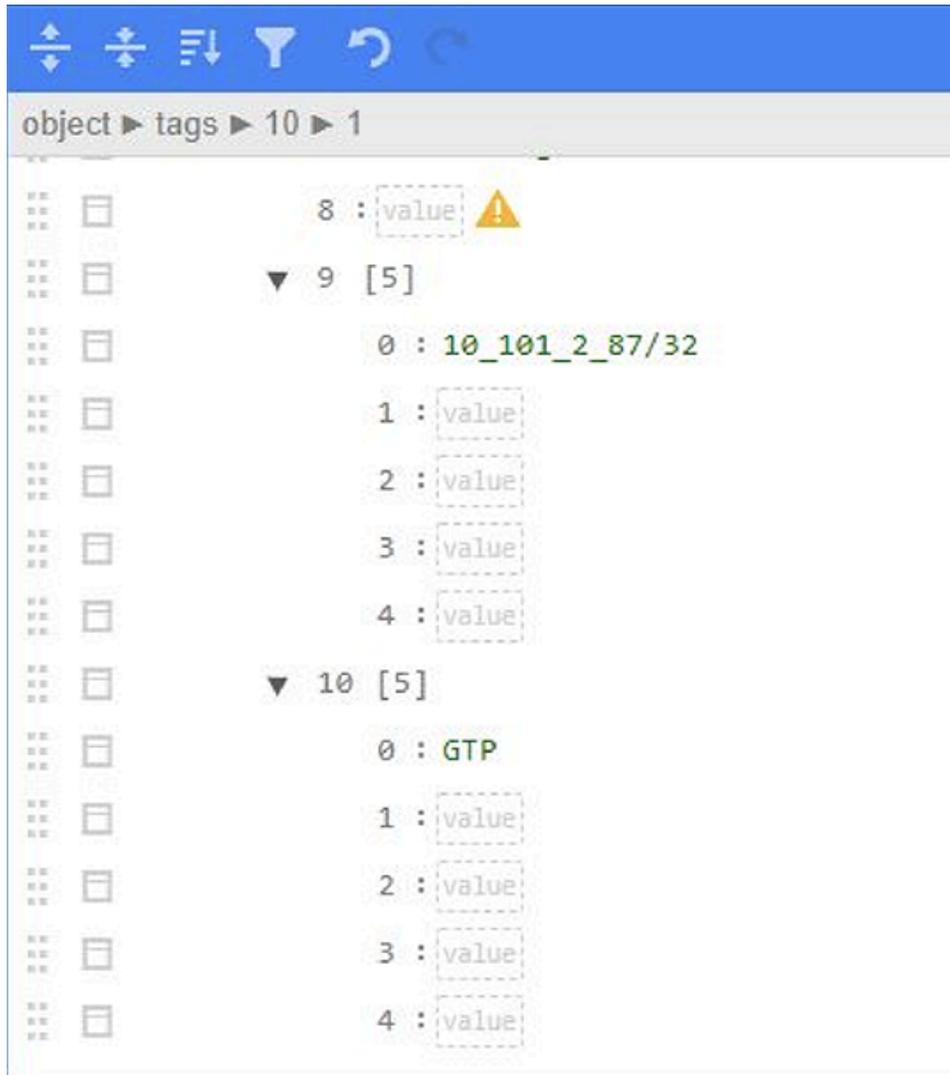
**Figure 7-10: Mapping a Tool to an IP Address: IP Block Edit**



2. To define a new IP block, append a range of IP addresses to the blocks section.

3. Scroll down and add a tag definition with the same number as the IP block.

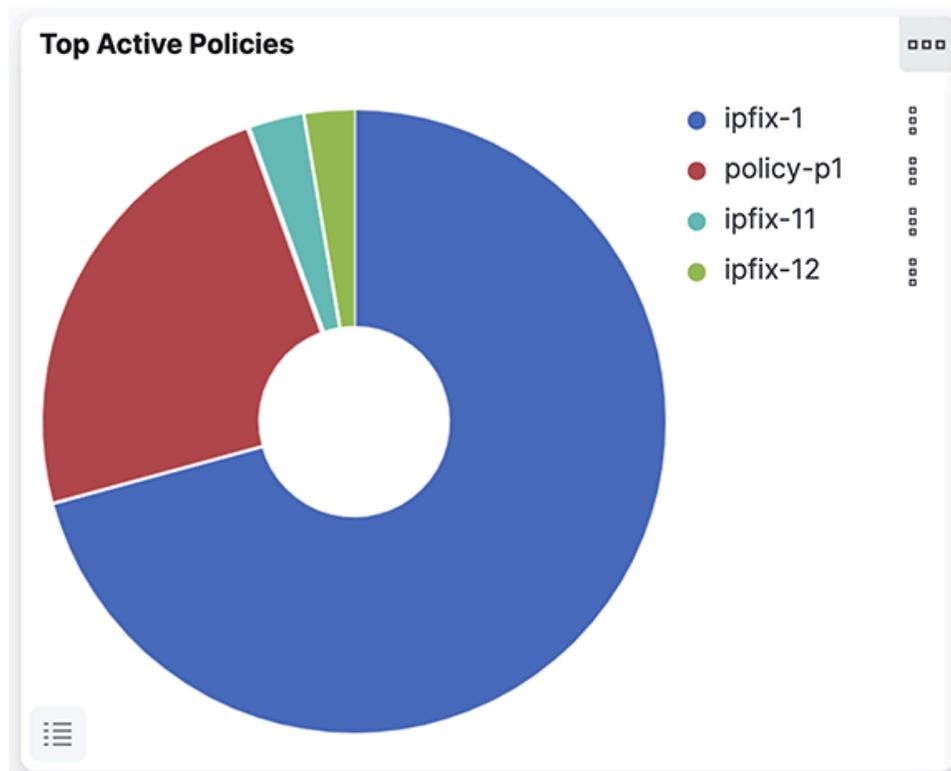
**Figure 7-11: Mapping a Tool to an IP Address: Define Tags**



4. Define the new IP block section tags, including a descriptive name for the specific tool.
5. Select **DMF Network > Policy Statistics**.

To cross-reference the information you get by labeling an IP block with information about any policies configured to forward traffic to that IP address.

**Figure 7-12: DMF Policies**



## 7.4 User Activity

This section identifies specific users transmitting or receiving traffic on the network.

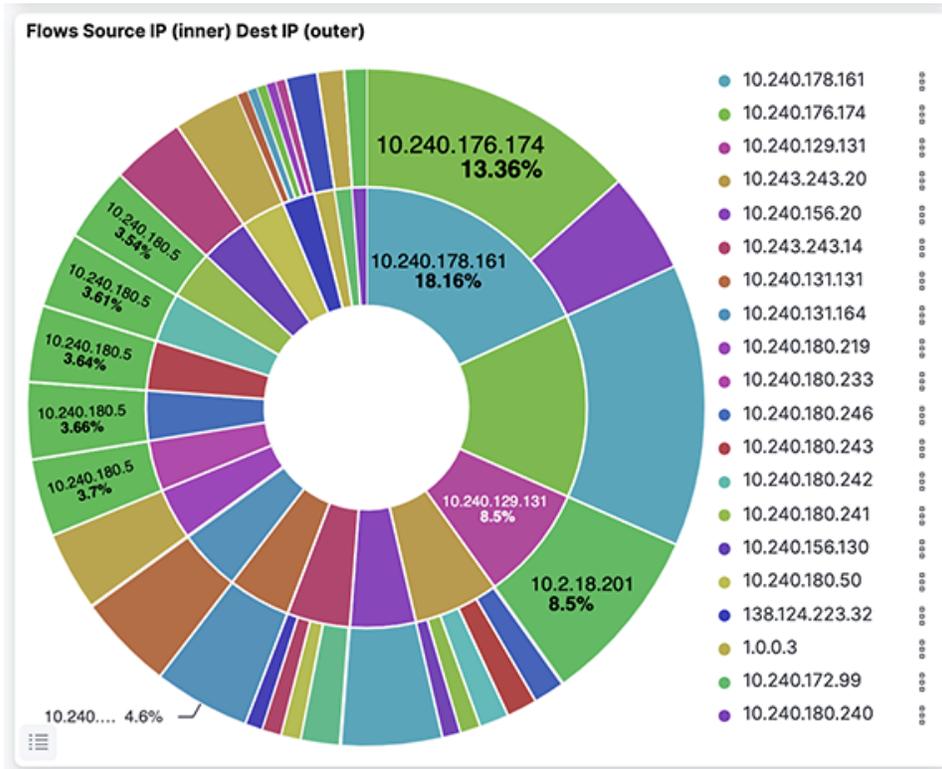
### 7.4.1 User Sessions

To identify users transmitting or receiving traffic on the network, use the following features:

- **Flows** dashboard
- **sFlow** dashboard
- **NetFlow** dashboard
- Open VPN or Active Directory mapping to IP address

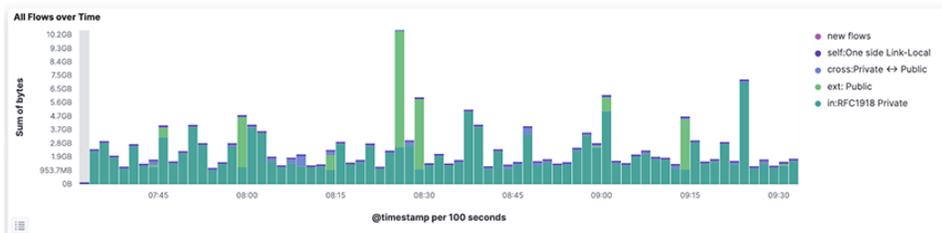
The **Flows** dashboards all provide an overall idea of who communicates on the network (traffic source and destination).

**Figure 7-13: Flows > Flows Source IP Dest IP**



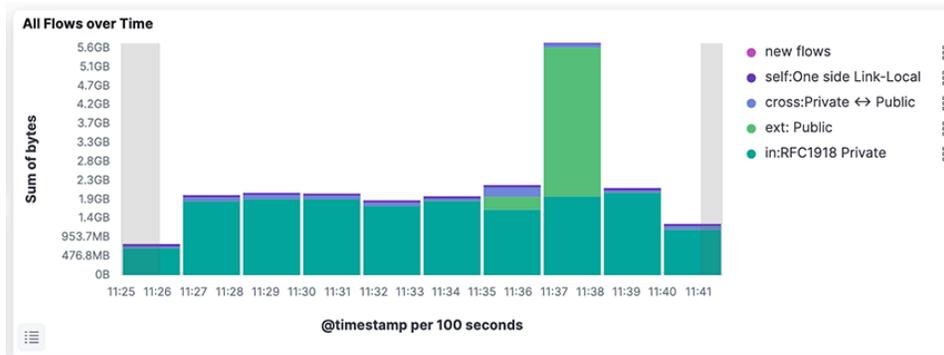
Click an IP address or L4 port, and with the + that appears, pin that to filter the dashboard for the selection. Every dashboard has a bar chart that shows traffic on the y-axis and time on the x-axis.

**Figure 7-14: All Flows Over Time**



To filter the display to a specific time, click and drag from left to right over the interesting period.

**Figure 7-15: All flows Over Time (Specific Time)**



It can also identify traffic associated with specific users after using the IP blocks configuration to map them to a specific IP address. Once saved, it can search for the user string to see traffic attributed to that user over the period displayed on the dashboard.

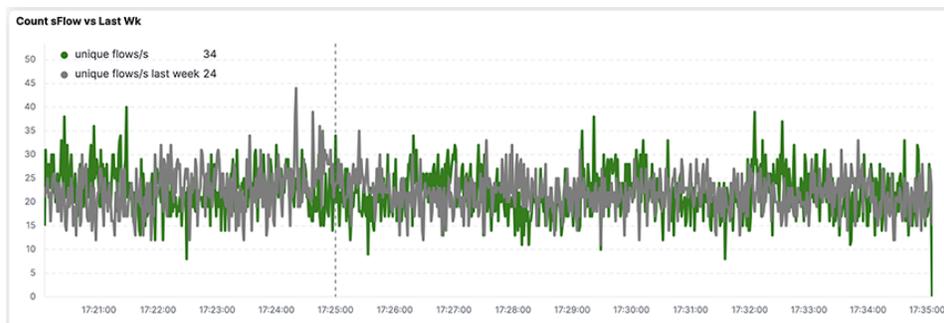
## 7.4.2 New Network Users

To identify new network users, use the following features:

- Comparing the same dashboard for two different periods
- **sFlow > Count sFlow vs Last Wk**
- **ARP** dashboard
- New Host Report

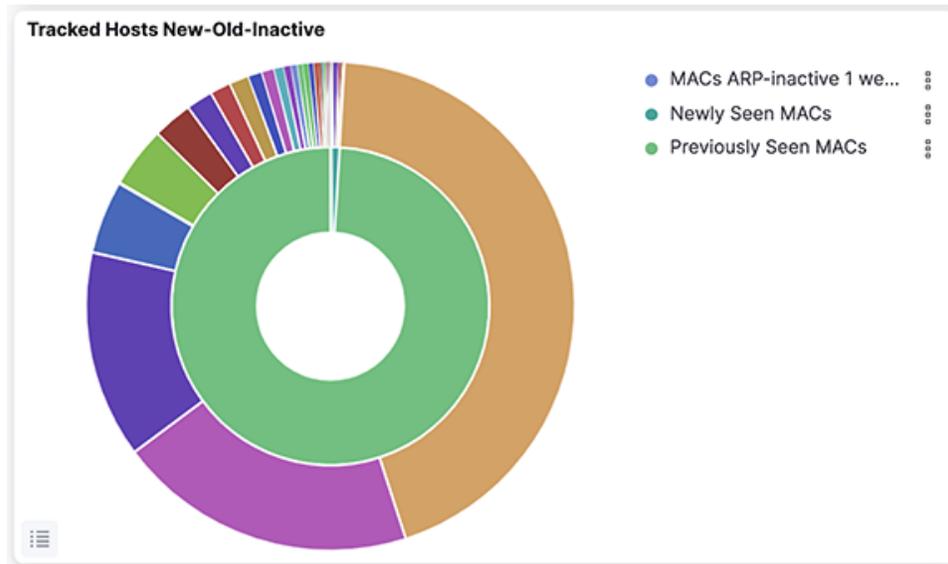
The sFlow dashboard provides a **Count sFlow vs Last Wk** visualization, which shows the number of unique flows being seen now vs. last week.

**Figure 7-16: sFlow > Count sFlow vs Last Wk**



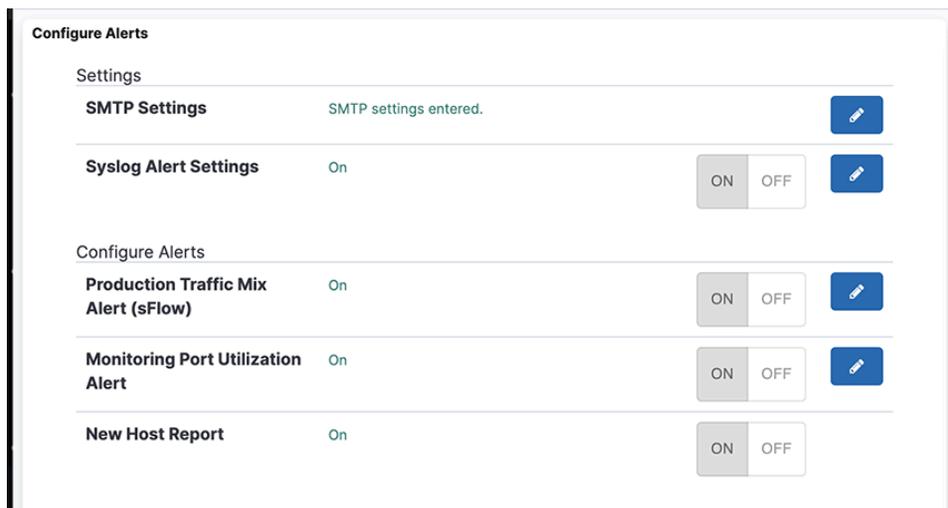
The **ARP** dashboard provides a visualization for **Tracked Hosts New-Old-Inactive, Vendor**.

**Figure 7-17: ARP > Tracked Hosts New-Old-Inactive, Vendor**



To use the **New Host** report, enable the report and configure where to send alerts on the **System > Configuration** page.

**Figure 7-18: System > Configuration > New Host Report**



### 7.4.3 Unauthorized Intranet Activity

To identify unauthorized usage of your internal network, use the following features:

- Malicious vs. compromised vs. apt zero-day vs. known threats. It enables the association of flows to users and flows to internal organizations.
- Searching by the username will reveal access to different organizations and their Apps.
- For OpenVPN users, when the IP is from a different geographical location, it shows the user's external IP. It may indicate a compromised account, especially in combination with access at odd hours.
- The OpenVPN server records logins with IP addresses and computer type, assigns IP addresses inside the lab, and sends Syslog on OpenVPN.

- Use the DMF Recorder Node to retrieve the original packets for forensic analysis and to obtain evidence of unauthorized activity.

## 7.5 Monitoring Active Directory Users

Windows Active Directory should be configured to audit logon and logoff events on Active Directory.

1. Download and install Winlogbeat from the Elastic website on the Windows machine. [Download Winlogbeat](#).
2. On the Analytics node, run: `sudo rm -rf *` inside `/home/admin/xcollector` and then run `docker exec xcollect /home/logstash/generate_client_keys.sh <AN IP> client`. It generates `.pem` files in `/home/admin/xcollector`.
3. On the Analytics node machine, replace the `winlogbeat.yml` file from `/opt/bigswitch/conf/x_collector/winlogbeat.yml` to the one in the Windows server. Edit the `logstash` output section:

```
#----- Logstash output -----
output.logstash:
#Point agent to analytics IPv4 in hosts below hosts: ["10.2.5.10:5043"]

#List of root certificates for HTTPS server verifications ssl.certificate_authorities: ["C:/Program Files/
Winlogbeat/security/ca/cacert.pem"]

#Certificate for SSL client authentication
ssl.certificate: "C:/Program Files/Winlogbeat/security/clientcert.pem"

#Client Certificate Key
ssl.key: "C:/Program Files/Winlogbeat/security/clientkey.pem"
```

4. Using the recovery account, use an SCP application to transfer the `.pem` files from the Analytics node to the Windows machine and update their locations in `winlogbeat.yml`.
5. On Windows, enter the powershell, navigate to `winlogbeat.exe`, and run: `.install-service-winlogbeat.ps1` to install **Winlogbeat**.
6. Test the configuration using `winlogbeat test config` to test `winlogbeat.yml` syntax and `winlogbeat test output` to test connectivity with `logstash` on the Analytics node.
7. Run `winlogbeat run -e` to start **Winlogbeat**.

# Monitoring Network Performance and Events

---

This chapter monitors network performance and identifies unusual events. It includes the following sections.

- [Interfaces Sending or Receiving Traffic](#)
- [Anomalies](#)
- [Application Data Management](#)
- [WAN Link Optimization](#)
- [Machine Learning](#)

## 8.1 Interfaces Sending or Receiving Traffic

To identify specific interfaces that are sending or receiving traffic, use the following features:

- DMF Top Filter interfaces

- Production interfaces

Figure 8-1: DMF Filter Interfaces

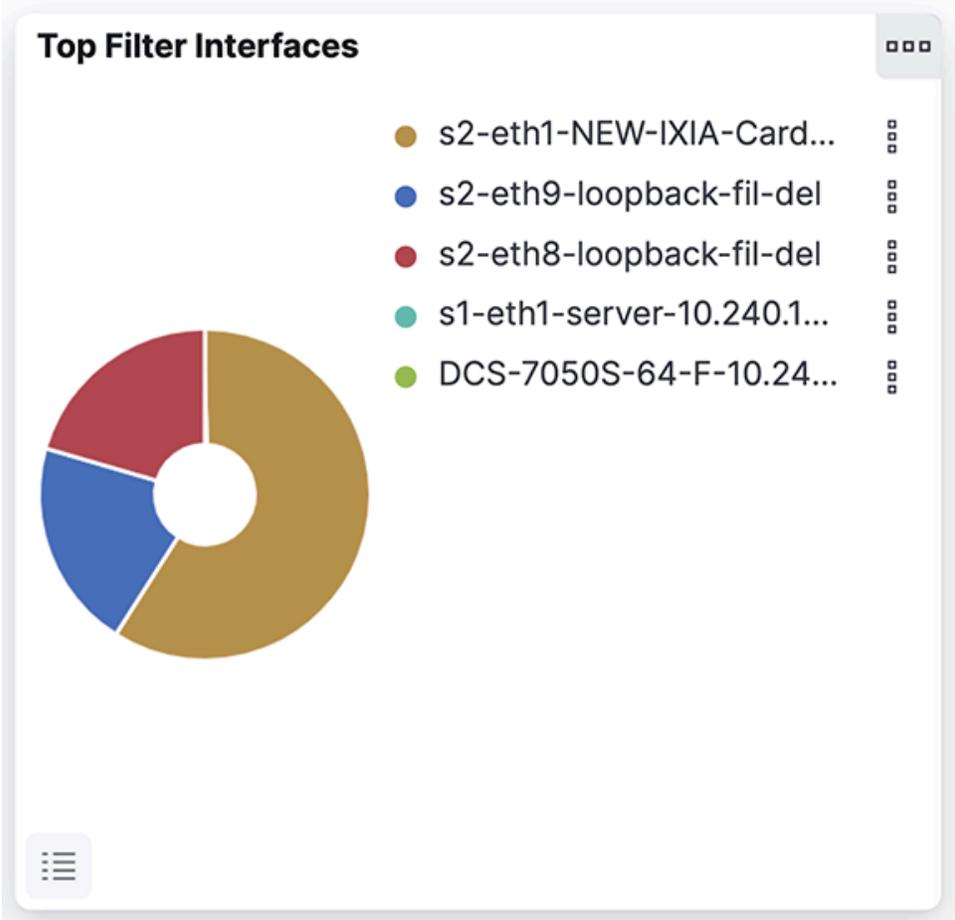
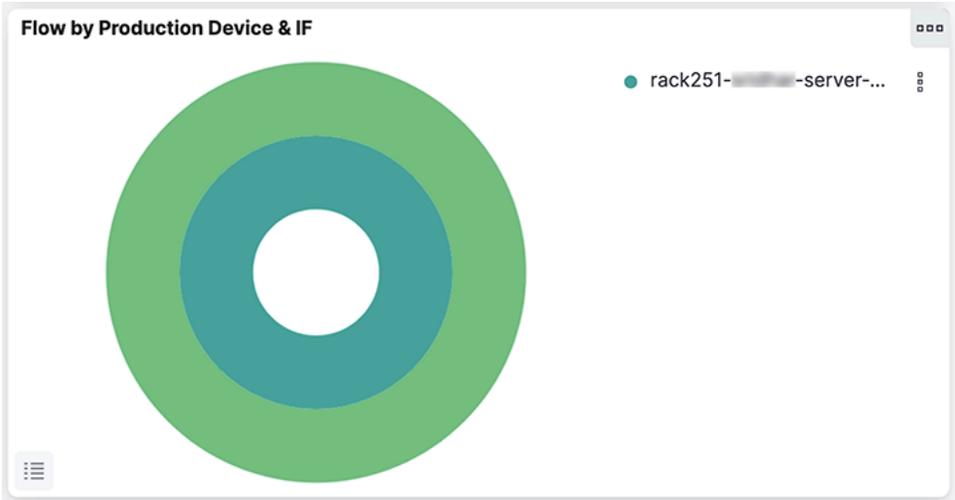


Figure 8-2: sFlow® > Flow by Production Device & IF



This information derives from the LLDP/CDP exchange between the production and DANZ Monitoring Fabric switches.

## 8.2 Anomalies

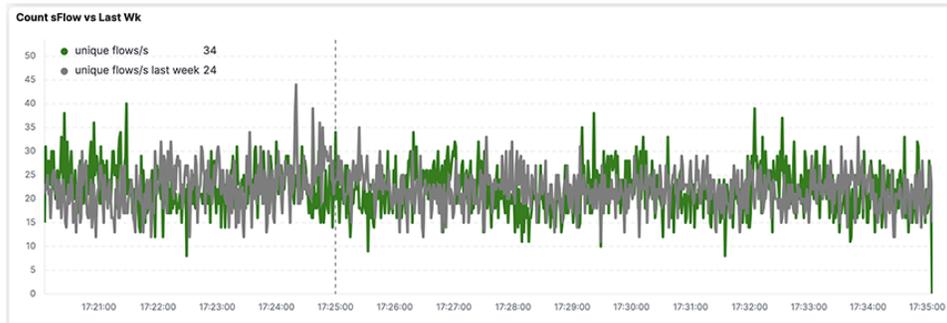
Use the following features to recognize unusual activity or events on the network.

- Comparing dashboards and visualization over time
- sFlow<sup>®\*</sup> > Count sFlow vs Last Wk
- New Flows & New Hosts
- Utilization alerts
- Machine Learning

Identify any unusual activity by comparing the same dashboard over the past 1 hour to the same time last week's data. For example, the bar visualization of traffic over time shows changing ratios of internal to external traffic, which can highlight an abnormality.

The **Count sFlow vs Last Wk** visualization in the **sFlow** dashboard shows the number of unique flows being seen now compared to last week. This visualization indicates unusual network activity and will help pinpoint a Denial of Service (DOS) attack.

**Figure 8-3: Count sFlow vs Last Wk**



In a well-inventoried environment, use the **New Flows & New Hosts** report.

**Figure 8-4: Production Traffic**

Configure Alerts

<b>Production Traffic Mix Alert (sFlow)</b>	Generates an alert when switch ports exceeds utilization threshold. Outbound Traffic Percentage <input type="text"/>	
<input type="button" value="Save"/> <input type="button" value="Cancel"/>		

Configure utilization alerts associated with the following DMF port types:

- Filter
- Delivery
- Core

\* sFlow<sup>®</sup> is a registered trademark of Inmon Corp.

- Services

Figure 8-5: Monitoring Port Utilization Alerts

**Monitoring Port Utilization Alert**

When this utilization exceeded send an alert.

All utilization (%)

Filter utilization (%)

Delivery utilization (%)

Core utilization (%)

Service utilization (%)

Managed Service utilization (%)



The other alerts available include the following.

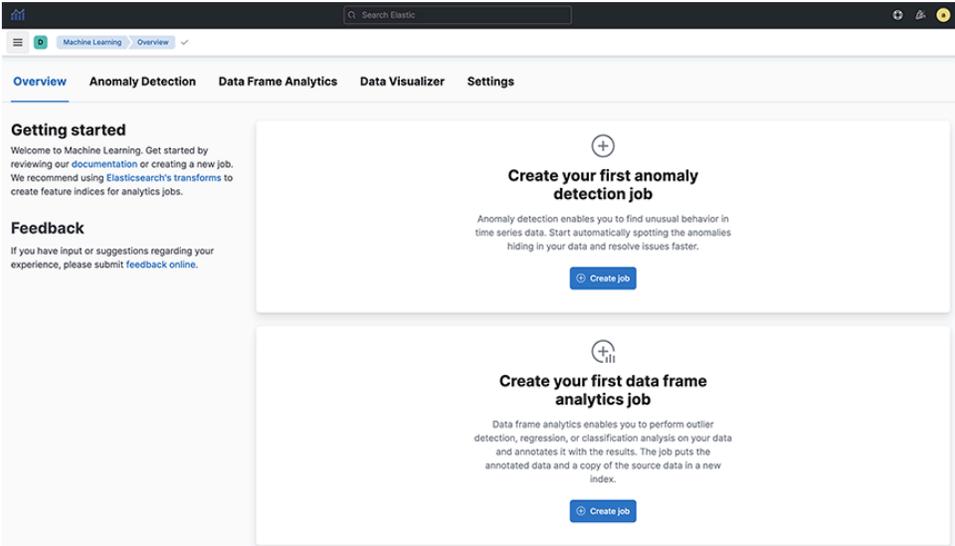
- The percentage of outbound traffic exceeds the usual thresholds.
- New hosts appear on the network every 24 hours.

Figure 8-6: New Host Report

**New Host Report** On

Perform Anomaly Detection in data over byte volume and characteristics over time using machine learning.

Figure 8-7: Machine Learning



### 8.3 Application Data Management

Application Data Management (ADM) helps users govern and manage data in business applications like SAP ERP. To use Arista Analytics for ADM, perform the following steps:

1. Pick a service IP address or block of IP addresses.
2. Identify the main body of expected communication with adjacent application servers.

3. Filter down to ports that need to be communicating.
4. Expand the time horizon to characterize necessary communication completely.
5. Save as CSV.
6. Convert the CSV to ACL rules to enforce in the network.

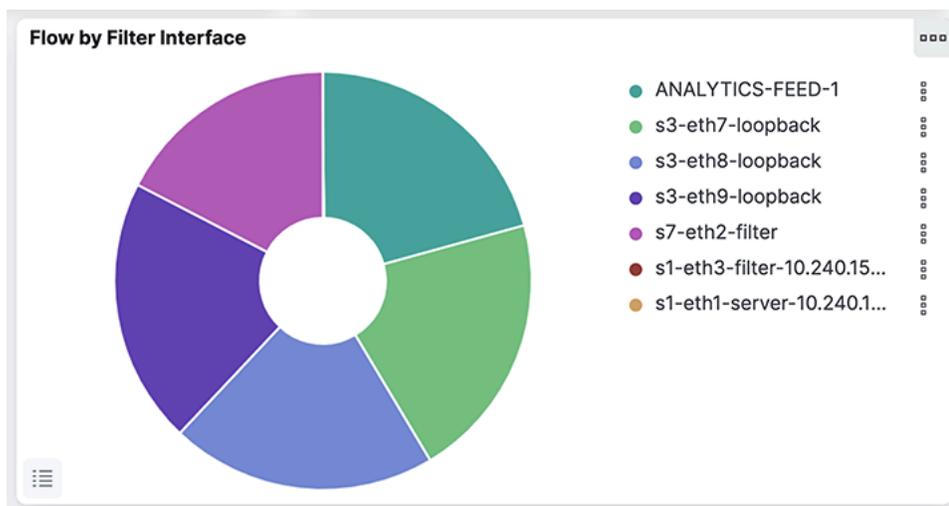
## 8.4 WAN Link Optimization

Use your knowledge of DMF filters or delivery interface names to monitor traffic to or from specific interfaces. DMF WAN interface names identified with a standard string, such as **wan**, can monitor the utilization of WAN links by reference to the DMF filter interface names.

To identify a WAN link or device that is approaching full utilization, complete the following steps:

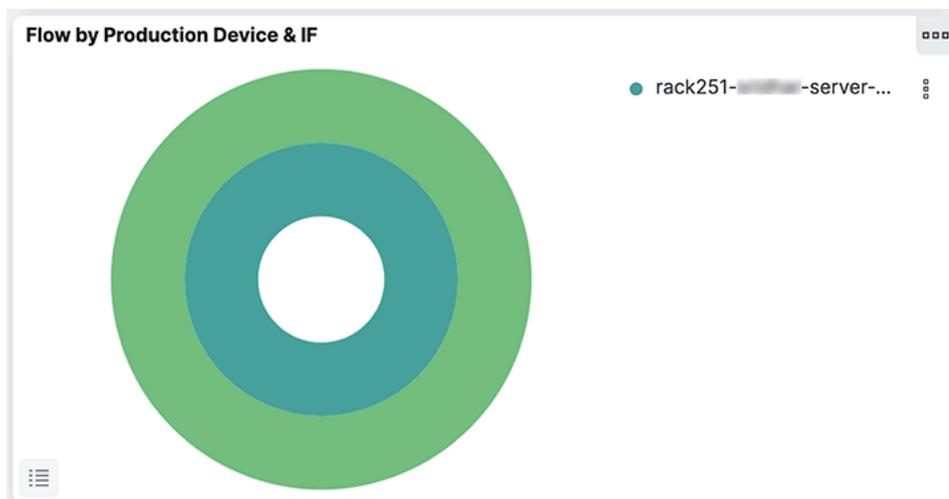
1. Select **sFlow**.
2. Refer to the **Flow by Filter Interface** visualization.

**Figure 8-8: Flow by Filter Interface**



This visualization displays the utilization for each DMF filter interface. To compare this to the traffic from the production interfaces (SPAN or Tap), use the **Flow by Production Device & IF** visualization.

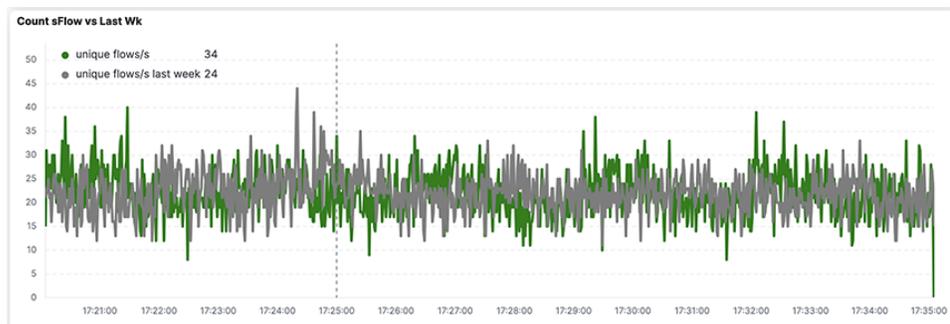
**Figure 8-9: Flow by Production Device & IF**



3. Select the Filter Interfaces corresponding to the WAN link.

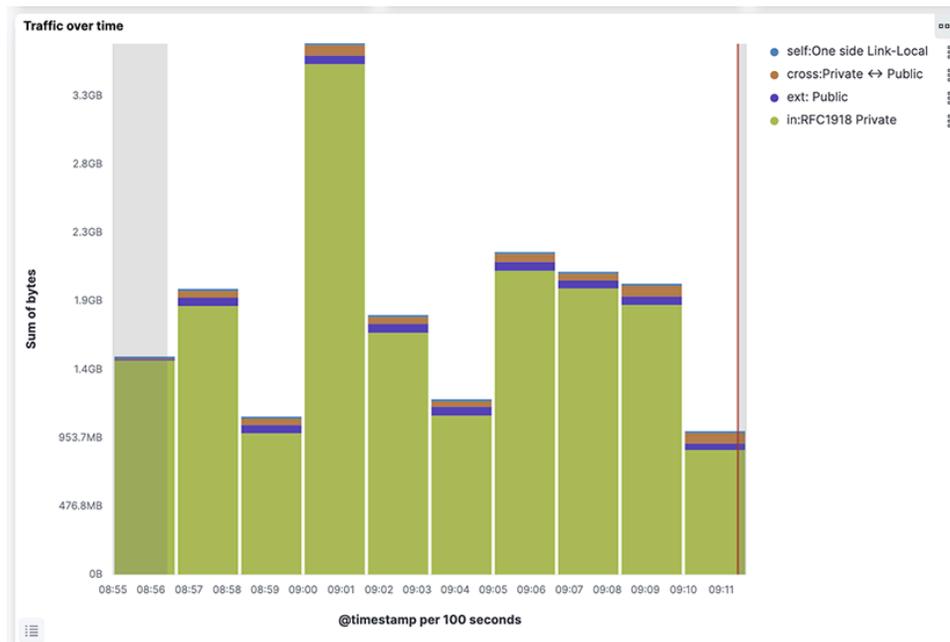
Refer to the **Count sFlow vs Last Wk** visualization to determine if any significant change in utilization has occurred.

**Figure 8-10: Count sFlow vs Last Wk**



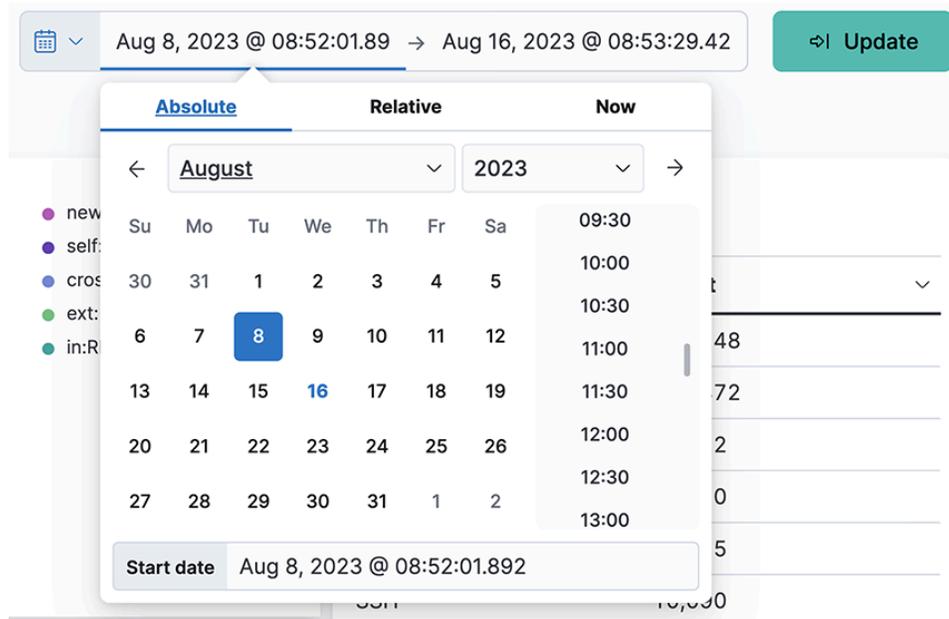
Use the **Traffic over Time** visualization to focus on peak and non-peak utilization periods. Drag the cursor horizontally over a peak utilization period, and the display is updated to zoom in on those events.

**Figure 8-11: Traffic Over Time**



4. Use the **Time Range** configuration to analyze traffic over a month for a more complete characterization.

**Figure 8-12: Expanding Time Period Using the Time Range**

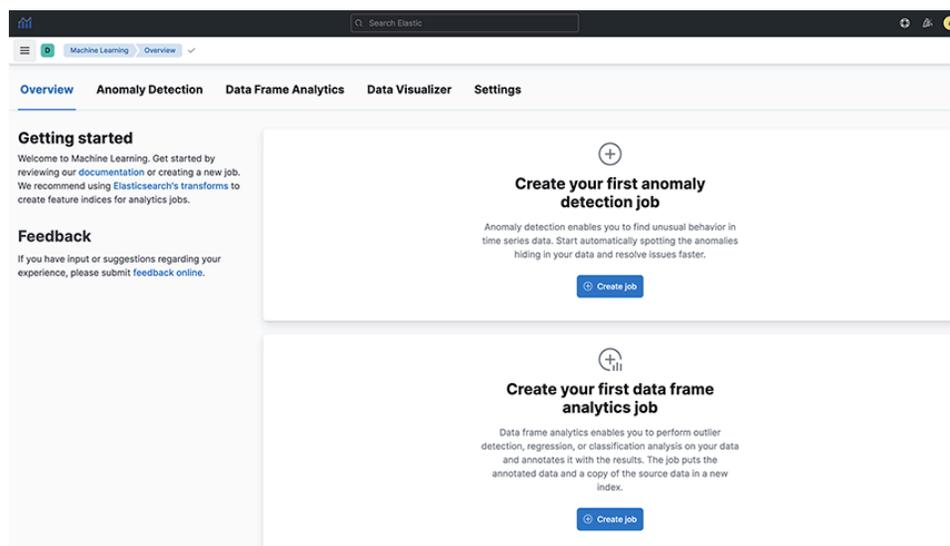


## 8.5 Machine Learning

Arista Analytics uses machine learning for anomaly detection. The following jobs are available:

- Single-metric anomaly detection
- Multimetric anomaly detection
- Population
- Advanced
- Categorization

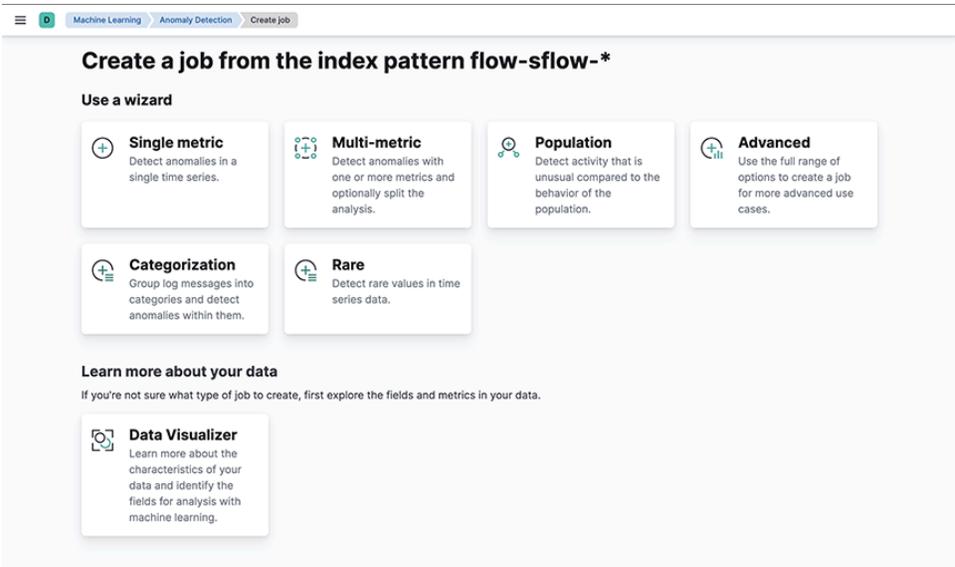
**Figure 8-13: Machine Learning**



For every job, a job ID must be configured. To create a machine learning job:

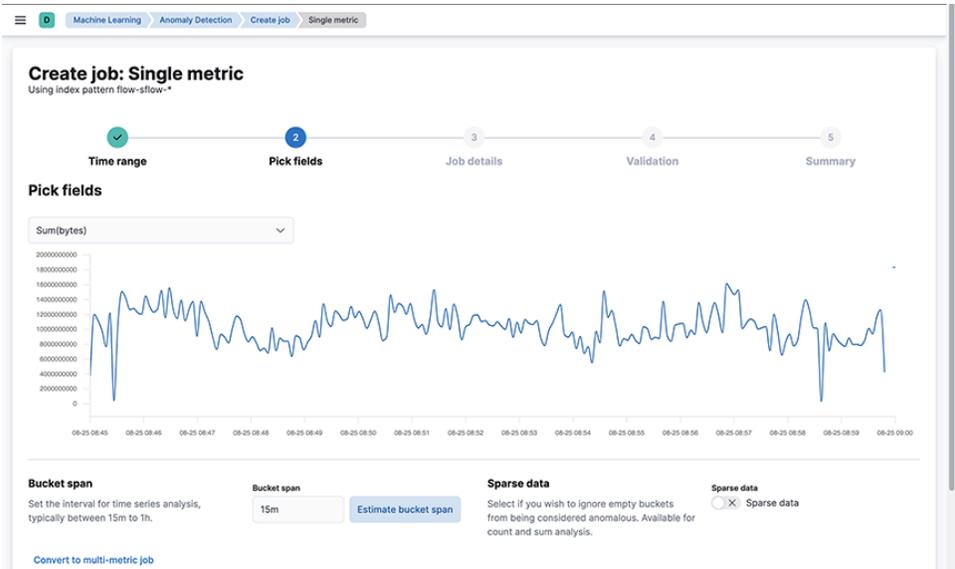
- Select the time range
- Select the appropriate metric
- Enter details: job ID, description, custom URLs, and calendars to exclude planned outages from the job

Figure 8-14: Machine Learning Job options



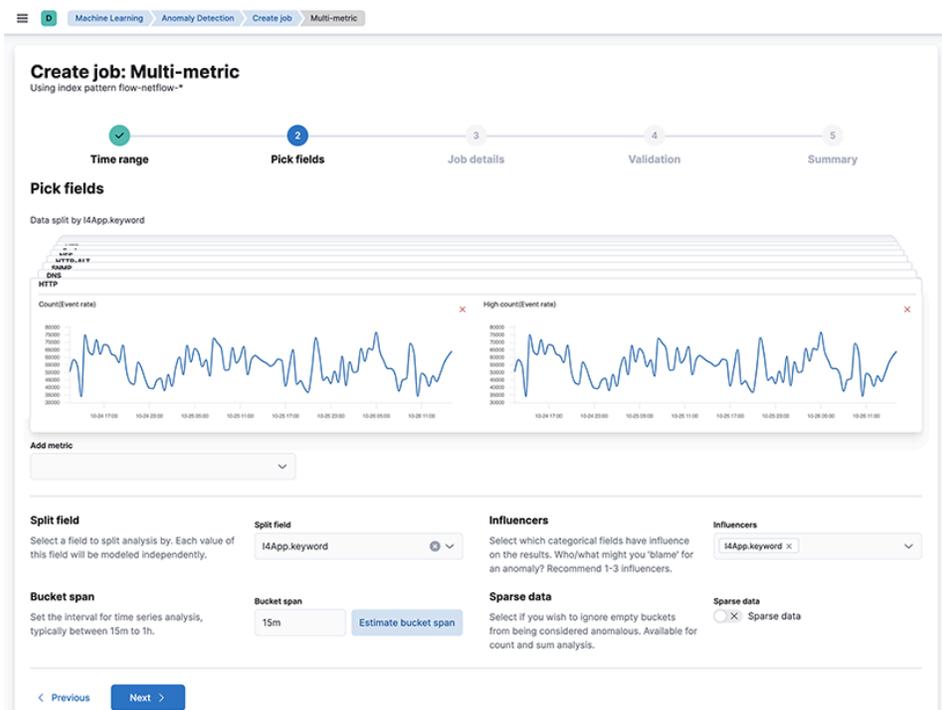
Single-metric anomaly detection uses machine learning on only one metric or field.

Figure 8-15: Single-metric Anomaly Detection



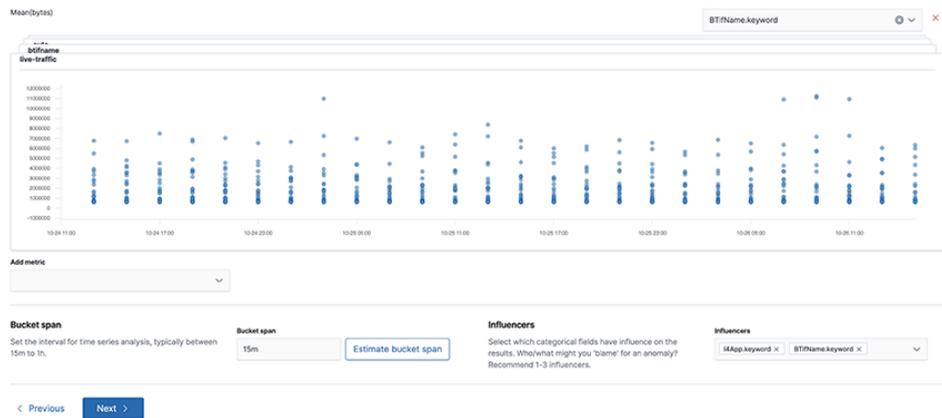
Multimetric and so on, I couldn't find any which anomaly detection uses machine learning on more than one metric field. The image below uses two metrics: over and running ml per L4 app.

**Figure 8-16: Multimetric Anomaly Detection**



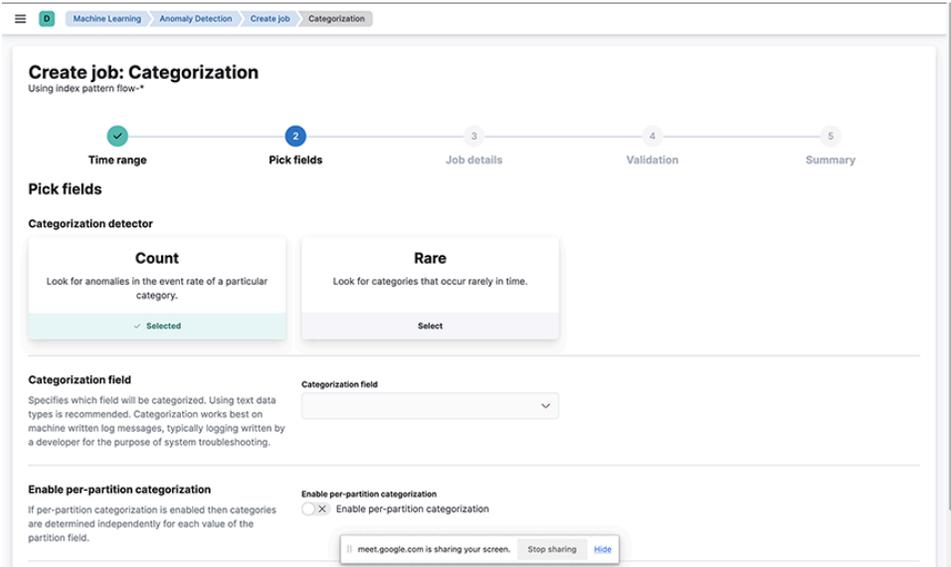
Multimetric Anomaly Detection detects network activity that differs from the population of data points. Arista Networks recommends this analysis for high-cardinality data.

**Figure 8-17: Population**



This job groups data points into categories and then finds anomalies between them.

Figure 8-18: Categorization



## Backup and Restore

This chapter includes the following sections.

- [Elasticsearch Snapshot and Restore](#)
- [Import and Export of Saved Objects](#)
- [Import and Export of Watchers](#)
- [Import and Export of Machine Learning Jobs](#)

### 9.1 Elasticsearch Snapshot and Restore

Elasticsearch provides a mechanism to snapshot data to a network-attached storage device and to restore from it.

1. Mount the Network File Storage (NFS) on the Analytics Node.
  - a. Create a directory on the remote Ubuntu Server (NFS store). This directory must have the user group **remoteuser** and **root**, respectively, with **10000** for the UID and **0** for the GID.
  - b. Stop the Elasticsearch container: `sudo docker elasticsearch stop`
  - c. Mount the remote store on `/opt/bigswitch/snapshot` in the Analytics server.
  - d. Start the Analytics Node: `sudo docker elasticsearch start`
2. Create a snapshot repository by running the following API call:

```
curl \
-k \
-X PUT \
-H 'Content-Type:application/json' \
-d '{"type":"fs","settings":{"location":"/usr/share/elasticsearch/sn
apshot"}}' \
-u admin:***** \
https://169.254.16.2:9201/_snapshot/test_automation
```

3. Take a snapshot by running the following API call:

```
curl \
-k \
-X POST \
-H 'Content-Type:application/json' \
-d '{"indices": ".ds-flow-sflow-stream-2023.08.21-000001", "include_glob
al_state": true, "ignore_unavailable": true, "include_hidden": true}' \
-u admin:***** \
https://169.254.16.2:9201/_snapshot/test_automation/test_snap1
```

4. To view the a snapshot, run the following API call:

```
curl \
-s -k \
-H 'Content-Type:application/json' \
-u admin:***** \
https://169.254.16.2:9201/_snapshot/test_automation/test_snap1?pretty
```

- To restore a snapshot, run the following API call:

```
curl \
-k \
-X POST \
-H 'Content-Type:application/json' \
-d '{ "indices": ".ds-flow-sflow-stream-2023.08.21-000001", "ignore_unavailable": true, "include_global_state": true, "rename_pattern": "(.+)", "rename_replacement": "restored_$1" }' \
-u admin:***** \
https://169.254.16.2:9201/_snapshot/test_automation/test_snap1/_restore
```

## 9.2 Import and Export of Saved Objects

The **Saved Objects** UI helps keep track of and manage saved objects. These objects store data for later use, including dashboards, visualization, searches, and more. This section explains the procedures for backing up and restoring saved objects in Arista Analytics.

### 9.2.1 Exporting Saved Objects

- Open the main menu, then click **Main Menu > Management > Saved Objects**.
- Select the custom-saved objects to export by clicking on their checkboxes.
- Click the **Export** button to download. Arista Networks suggests changing the file name to the nomenclature that suits your environment (for example, `clustername_date_saved_objects_<specific_name_or_group_name>.ndjson`).



**Note:** Arista Networks recommends switching ON **include related objects** before selecting the **export** button. If there are any missing dependency objects, selecting **include related objects** may throw errors, in which case switch it OFF.

- The system displays the following notification if the download is successful.

**Figure 9-1: Verifying a Saved/Downloaded Object**

The screenshot shows the Arista Analytics interface. The main content area is titled "Saved Objects" and contains a table of objects. The table has columns for "Type", "Title", "Tags", and "Actions". Several objects are selected with checkboxes. A notification at the bottom right of the screen states: "Your file is downloading in the background".

Type	Title	Tags	Actions
<input checked="" type="checkbox"/>	Flow Selectors		...
<input checked="" type="checkbox"/>	Flow Heatmap VPCs		...
<input type="checkbox"/>	ICMP Messages Over Time		...
<input checked="" type="checkbox"/>	ICMP By Error Description		...
<input checked="" type="checkbox"/>	ICMP Top Err Dest IPs		...
<input checked="" type="checkbox"/>	ICMP Error Dest Port Apps		...
<input type="checkbox"/>	ICMP by Filter Interface		...
<input type="checkbox"/>	ICMP by Production Device & IF		...
<input checked="" type="checkbox"/>	Top ICMP Message Source		...
<input checked="" type="checkbox"/>	Top ICMP Message Dest		...
<input checked="" type="checkbox"/>	NF over Time		...
<input type="checkbox"/>	NF by QoS PHB		...
<input type="checkbox"/>	NF Top Talkers by Flow (Table)		...
<input type="checkbox"/>	nFlow Live L4 Ports		...
<input type="checkbox"/>	nFlow by Filter Interface		...



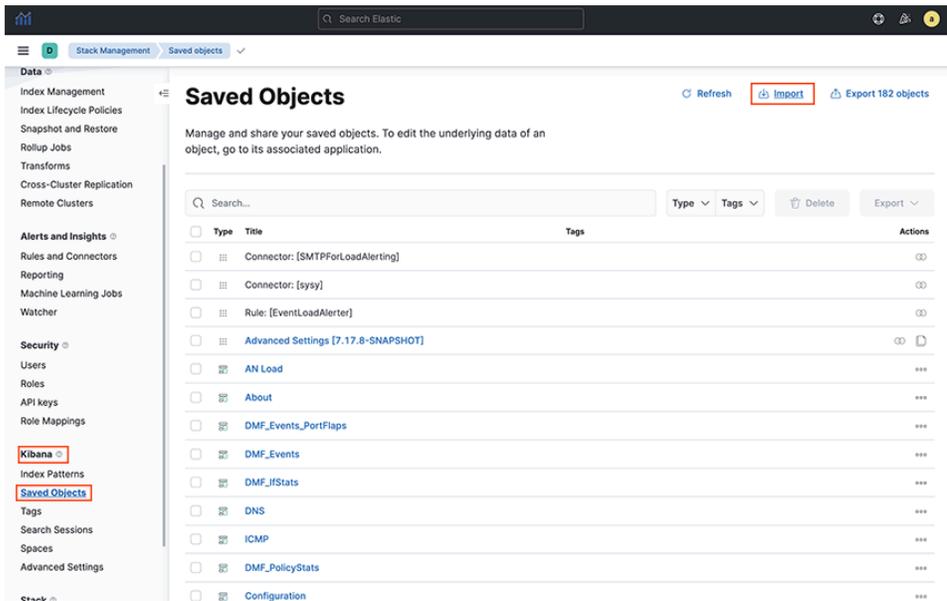
### Note: Recommended Best Practices

- While creating saved objects, Arista Networks recommends naming conventions that suit your environment. For instance, in the example above, a naming pattern has been used, prefixed with “ARISTA” and specifying **Type: dashboard**, which allows a manageable set of items to click individually or to select all. Furthermore, exporting individual dashboards based on their **Type** is a more appropriate option, as tracking modifications to a dashboard improves using this method. Dashboards should use only custom visualizations and searches (i.e., do not depend on default objects that might change during a software upgrade).
- Do not edit any default objects. Arista Networks suggests saving the new version with a different (custom) name if default objects require editing.
- The files exported should be treated as code and reserved in a source control system, so dissimilarities and rollbacks are possible under standard DevOps approaches.

## 9.2.2 Importing Saved Objects

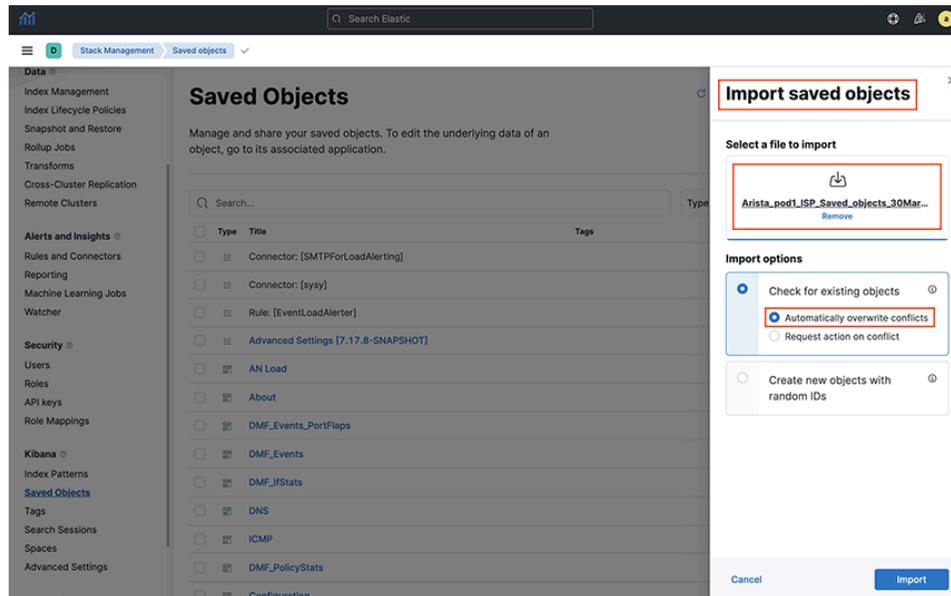
1. To import one or a group of custom-created objects, navigate to **Main Menu > Management > Kibana > Saved Objects**.

Figure 9-2: Importing a Group of Saved Objects



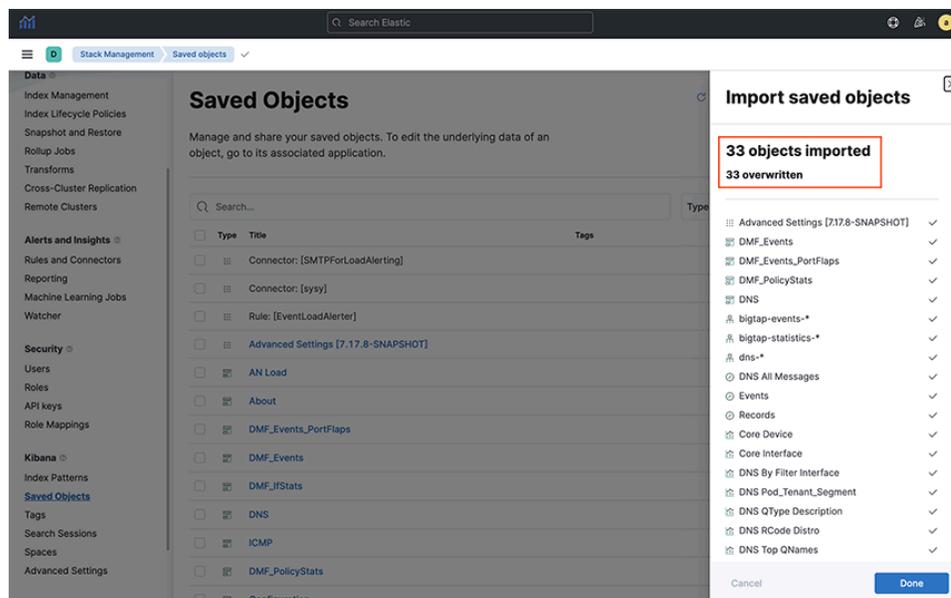
- Click **Import** and navigate to the NDJSON file that represents the objects to import. By default, saved objects already in Kibana are overwritten by the imported object. The system should display the following screen.

**Figure 9-3: NDJSON File Import Mechanism**



- Verify the number of successfully imported objects. Also verify the list of objects, selecting **Main Menu > Management > Kibana > Saved Objects > search for imported objects**.

**Figure 9-4: Import Successful Dialog Box**



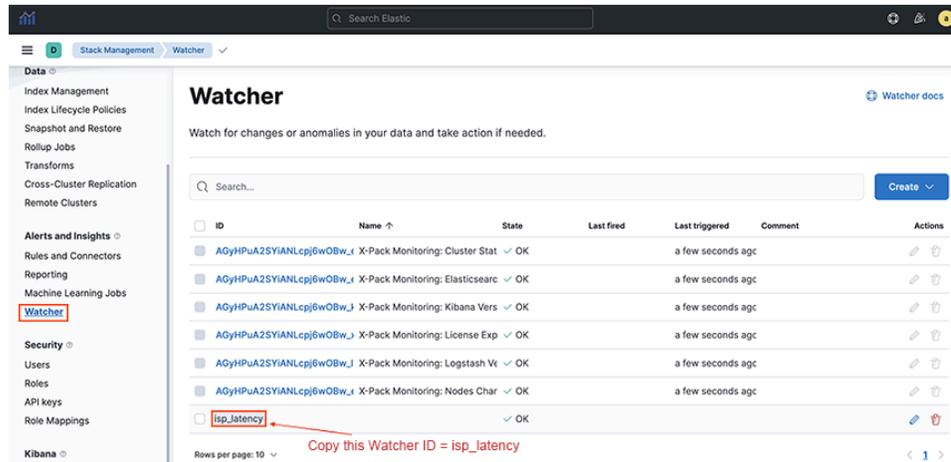
## 9.3 Import and Export of Watchers

Use the Watcher feature to create actions and alerts based on certain conditions and periodically evaluate them using queries on the data. This section explains the procedure of backing up and restoring the Watchers in Arista Analytics.

## 9.3.1 Exporting Watchers

1. The path parameter required to back up the Watcher configuration is `watcher_id`. To obtain the `watcher_id`, go to **Main Menu > Management > Watcher > Watcher\_ID**.

Figure 9-5: Find Watcher\_ID



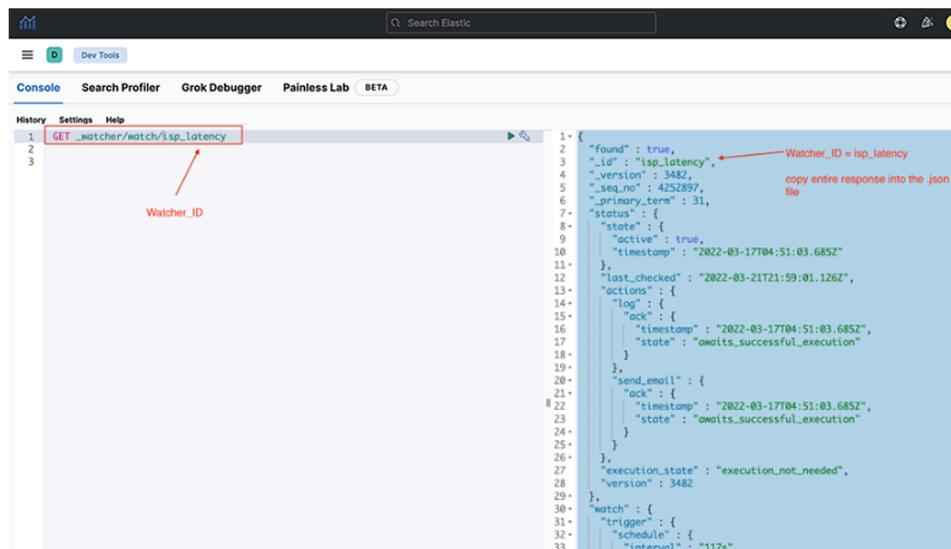
2. Open the main menu, then select **Dev Tools > Console**. Issue the **GET** API mentioned below with the `watcher_id`. The response appears in the output terminal.

Run the following API call:

```
GET _watcher/watch/<watcher_id>
```

Replace **Watcher\_ID** with the `watcher_id` name copied in **Step 1**.

Figure 9-6: GET API

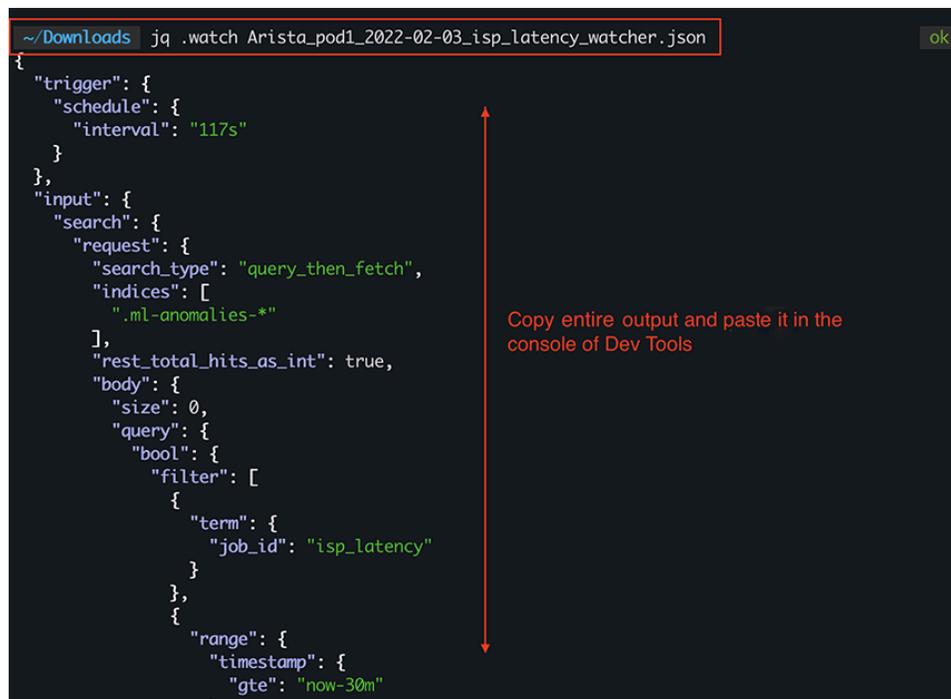


3. Copy the API response from **Step 2** into a `.json` file with the terminology that suits the environment, and keep track of it. As an example, the following may be helpful to nomenclature: `Arista_pod1_2022-02-03_isp_latency_watcher.json`.

## 9.3.2 Importing Watchers

1. Not all exported fields are needed when importing a Watcher. To filter out the unwanted fields from the exported file, use the `jq` utility. Use `jq .watch <exported_watcher.json>` and import the output.

Figure 9-7: jq Command Output

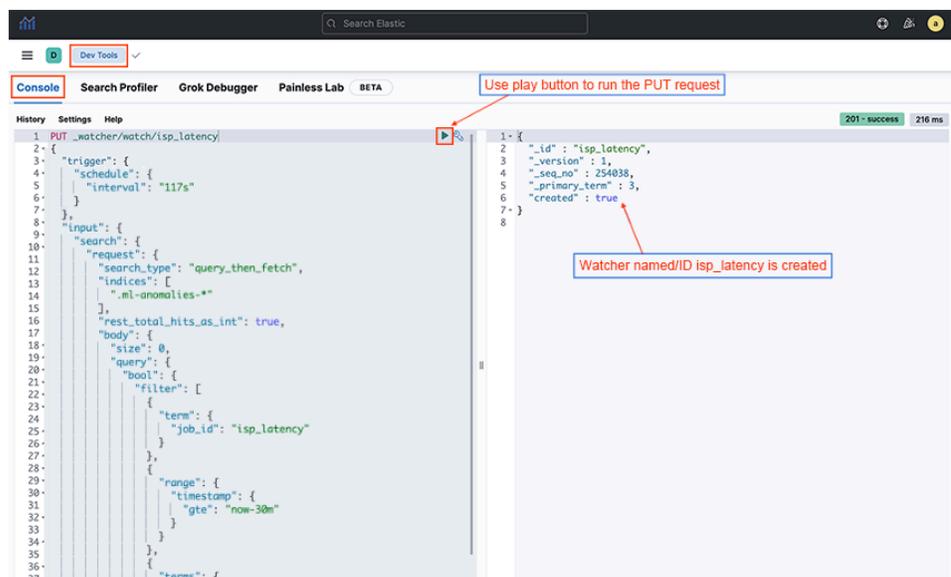


```
~/Downloads jq .watch Arista_pod1_2022-02-03_isp_latency_watcher.json ok
{
  "trigger": {
    "schedule": {
      "interval": "117s"
    }
  },
  "input": {
    "search": {
      "request": {
        "search_type": "query_then_fetch",
        "indices": [
          ".ml-anomalies-*"
        ],
        "rest_total_hits_as_int": true,
        "body": {
          "size": 0,
          "query": {
            "bool": {
              "filter": [
                {
                  "term": {
                    "job_id": "isp_latency"
                  }
                },
                {
                  "range": {
                    "timestamp": {
                      "gte": "now-30m"
                    }
                  }
                }
              ]
            }
          }
        }
      }
    }
  }
}
```

Copy entire output and paste it in the console of Dev Tools

2. Click **DevTools** > **console**, enter the API `PUT _watcher/watch/<watcher_id>`, and copy the **Step 1** output into the following screen. Replace **Watcher\_ID** with the desired Watcher name. The output terminal will confirm the creation of the Watcher.

Figure 9-8: PUT API in Dev Tools Console



Use play button to run the PUT request

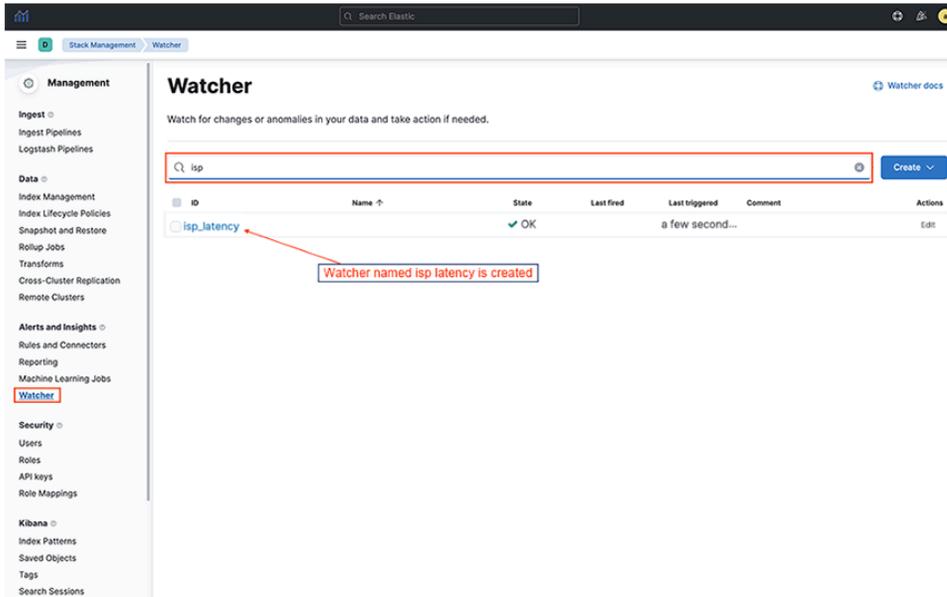
```
1 PUT _watcher/watch/isp_latency
2 {
3   "trigger": {
4     "schedule": {
5       "interval": "117s"
6     }
7   },
8   "input": {
9     "search": {
10      "request": {
11        "search_type": "query_then_fetch",
12        "indices": [
13          ".ml-anomalies-*"
14        ],
15        "rest_total_hits_as_int": true,
16        "body": {
17          "size": 0,
18          "query": {
19            "bool": {
20              "filter": [
21                {
22                  "term": {
23                    "job_id": "isp_latency"
24                  }
25                },
26                {
27                  "range": {
28                    "timestamp": {
29                      "gte": "now-30m"
30                    }
31                  }
32                }
33              ]
34            }
35          }
36        }
37      }
38    }
39  }
40 }
```

```
1- {
2   "_id": "isp_latency",
3   "_version": 1,
4   "_seq_no": 254038,
5   "_primary_term": 3,
6   "created": true
7 }
8
```

Watcher named/ID isp\_latency is created

3. Locate the newly created Watcher in the **Main menu > Management > Elasticsearch > Watcher > search with Watcher\_ID**.

**Figure 9-9: Watcher**



## 9.4 Import and Export of Machine Learning Jobs

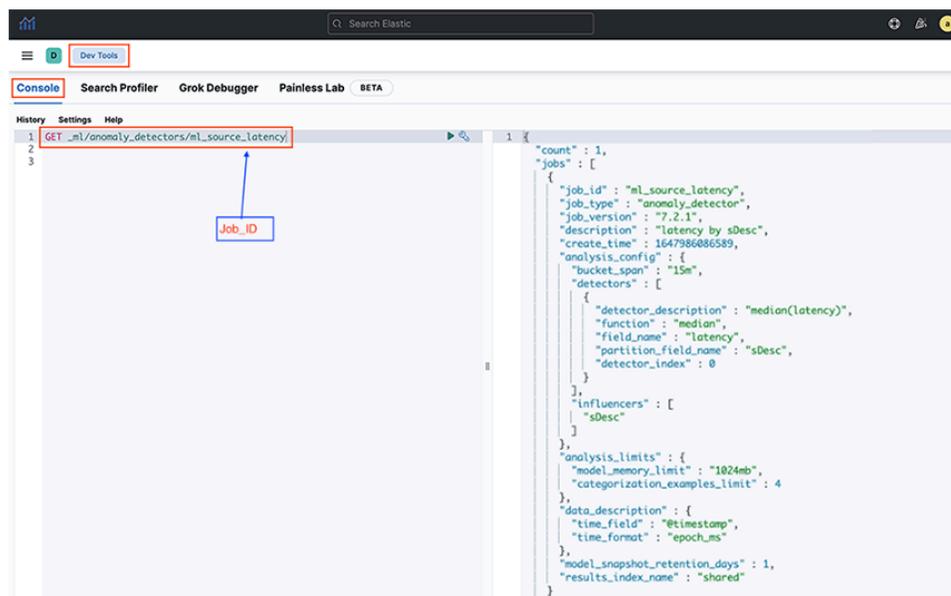
Machine Learning (ML) automates time series data analysis by creating accurate baselines of normal behavior and identifying anomalous patterns. This section explains ways to back up and restore the Machine Learning jobs in Arista Analytics.

### 9.4.1 Exporting Machine Learning Jobs

1. Open the main menu, then select **Dev Tools > Console**. Send a **GET `_ml/anomaly_detectors/<Job-id>`** request to Elasticsearch and view the response of all the Machine Learning anomaly jobs.

Replace **Job\_id** with the ML job name. The system displays the following output when executing the **GET** request.

**Figure 9-10: Main Menu > Dev Tools > Console**



2. Copy the **GET** API response of the ML job into a `.json` file with terminology that suits your environment and keep track of it. An example of appropriate nomenclature might be `Arista_pod1_2022-02-03_ML_Source_Latency_ML_job.json`.

## 9.4.2 Importing Machine Learning Jobs

1. It is optional to import all the exported fields. Only **description**, **analysis\_config**, and **data\_description** fields may be needed. Running `jq '.jobs[] | {description, analysis_config, data_description}' <json-filename>` copies the output into the **Dev tools** console. Replace **json-filename** with the filename of the JSON file previously exported.

Run the following API call:

```
jq '.jobs[] | {description, analysis_config, data_description}' Arista_pod1_2022-02-03_ML_Source_Latency_ML_job.json
```

**Figure 9-11: jq Required Fields**



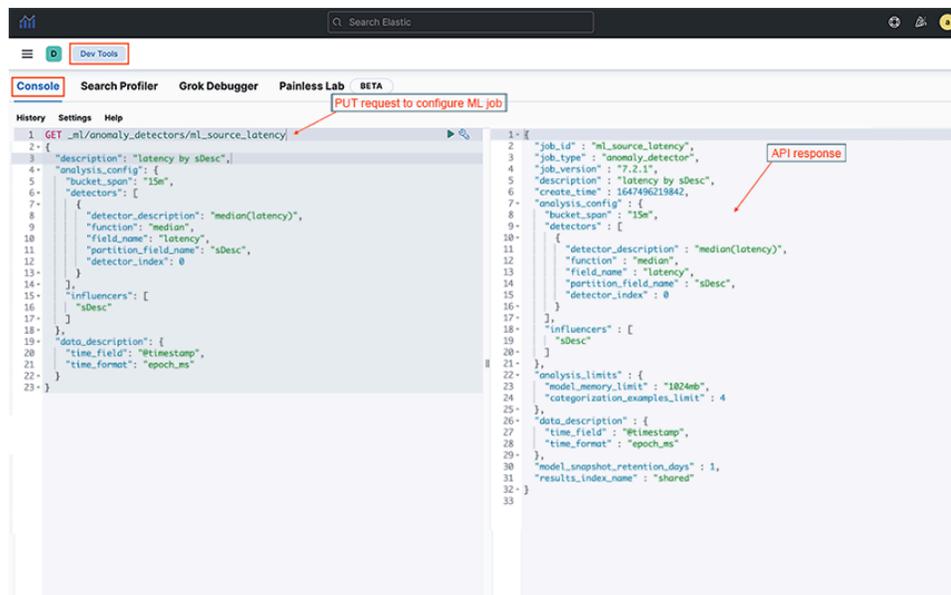
2. Select **Dev tools > Console** and copy the **Step 1** output into the screen below and the **PUT** request.

Run the following API call:

```
PUT _ml/anomaly_detectors/<ml_job_name>
```

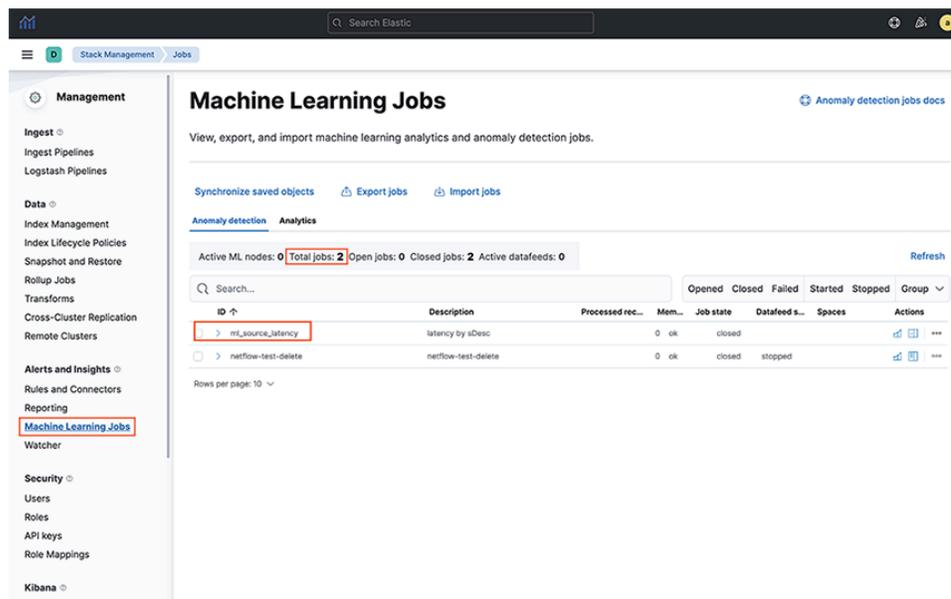
Replace *ml\_job\_name* with the specific string of the ML Job name.

Figure 9-12: PUT ML Jobs API



3. The successful response to the **PUT** request confirms the creation of the ML Job. Further, verify imported ML jobs by selecting **Main menu > Machine Learning > Job Management > search with ML Job Name**.

Figure 9-13: ML Job Verification



## Using TACACS+ and RADIUS to Control Access to the Arista Analytics CLI

This appendix describes using TACACS+ and RADIUS servers to control administrative access to the Analytics Node.

### 10.1 Using AAA Services with Arista Analytics

Use remote Authentication, Authorization, and Accounting (AAA) services using TACACS+ or RADIUS servers to control administrative access to the Analytics Node CLI.

The following table lists the accepted Attribute-Value (AV) pairs:

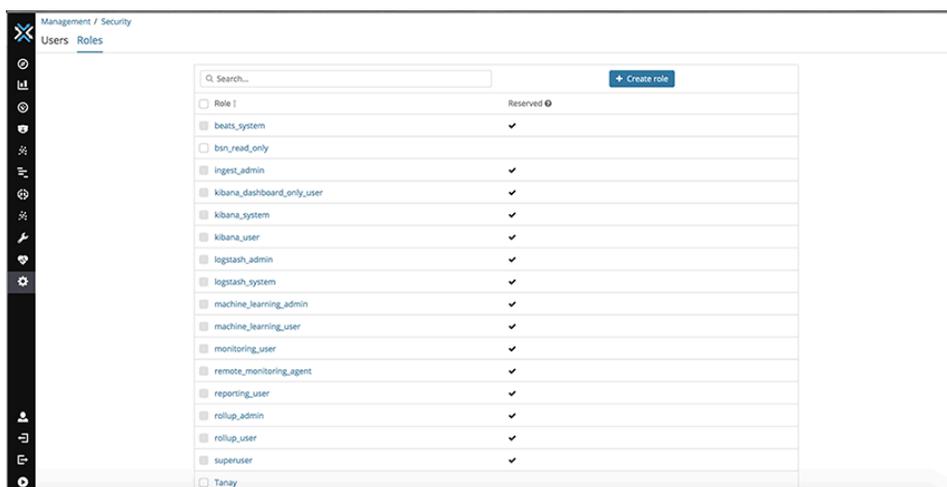
Attributes	Values
BSN-User-Role	admin read-only bigtap-admin bigtap-read-only



**Note:** The remotely authenticated **admin** and **bigtap-admin** users and the **read-only** and **bigtap-read-only** users have the same privileges. The **bigtap-admin** and **bigtap-read-only** values are supported to create BMF-specific entries without affecting the **admin** and **read-only** TACACS+ server entries.

You must also create a role in Elasticsearch with the same name as the group configured in the CLI.

**Figure 10-1: Creating a Group in Elasticsearch**



A remotely authenticated admin user has full administrative privileges. Read-only users on the switch must be remotely authenticated. Read-only access is not configurable for locally authenticated user accounts.

---

Read-only users can only access login mode, from which they can view most **show** commands, with some limitations, including the following:

- TACACS, SNMP, and user configuration are not visible to the **read-only** user in the output from the **show running-config** command.
- **show snmp**, **show user**, and **show support** commands are disabled for the **read-only** user.



**Note:** Local authentication and authorization take precedence over remote authentication and authorization.

Privileges at the remote TACACS+ server must be configured using the following attribute-value pairs:

- **Supported attribute name:** BSN-User-Role
- **Supported attribute values:** admin, read-only

Use a TACACS+ server to maintain administrative access control instead of using the Analytics Node local database, however, it is a best practice to keep the local database as the secondary authentication and authorization method in case the remote server becomes unavailable.

## 10.1.1 DMF TACACS+ Configuration

The DANZ Monitoring Fabric (DMF) requires the following configuration on TACACS+ servers and the configuration required on the Analytics Node.

### Authentication Method

- Configure the TACACS+ server to accept ASCII authentication packets. Do not use the **single connect only** protocol feature.
- The DMF TACACS+ client uses the ASCII authentication method. It does not use PAP.

### Device Administration

- Configure the TACACS+ server to connect to the device administration login service.
- Do not use a network access connection method, such as PPP.

### Group Memberships

- Create a **bigtap-admin** group. Make all DANZ Monitoring Fabric users part of this group.
- TACACS+ group membership is specified using the BSN-User-Role AV Pair as part of TACACS+ session authorization.
- Configure the TACACS+ server for session authorization, not for command authorization.



**Note:** The BSN-User-Role attribute must be specified as **Optional** in the **tac\_plus.conf** file to use the same user credentials to access ANET and non-ANET devices.

### Enabling Remote Authentication and Authorization on the Analytics Node

Use the following commands to configure remote login authentication and authorization. The examples use the SSH default for connection type.

```
analytics-1# tacacs server host 10.2.3.201
analytics -1# aaa authentication login default group tacacs+ local
analytics -1# aaa authorization exec default group tacacs+ local
```

All users in the **bigtap-admin** group on TACACS+ server **10.2.3.201** have full access to the Arista Analytics Node.

## User Lockout

Use the following command to lock out an AAA user after a calculated number of incorrect login attempts.

```
(config)#aaa authentication policy lockout failure F window W duration D
max-failures = F = [1..255] duration = D = [1..(2^32 - 1)] window = W = [1..
(2^32 - 1)]
```

## 10.2 Adding a TACACS+ Server

To view the current TACACS+ configuration, enter the `show running-config` command, as in the following example:

```
analytics -1(config-switch)# show run switch BMF-DELIVERY-SWITCH-1 tacacs
override-enabled
tacacs server host 1.1.1.1 key 7 020700560208
tacacs server key 7 020700560208
analytics -1(config-switch)#
```

It displays the TACACS+ key value as a type7 secret instead of plaintext.

Complete the following steps to configure the Analytics Node with TACACS+ to control administrative access to the switch.

Identify the IP address of the TACACS+ server and any key required for access using the `tacacs server` command, which has the following syntax:

```
tacacs server <server> [key {<plaintext-key> | 0 <plaintext-key> | 7
<encrypted-key>}]
```

You can enable up to four AAA servers by repeating this command for each server. For example, using a plaintext key, the following command enables TACACS+ with the server running at **10.2.3.4**.

```
analytics -1(config-switch)# tacacs server 10.1.1.1 key 0 secret
```

In case of a missing key, it uses an empty key.



**Note:** Do not use the pound character (#) in the TACACS secret. It is the start of a comment in the PAM config file.

Each TACACS+ server connection can be encrypted using a pre-shared key.

To specify a key for a specific host, use one of the following commands:

```
analytics -1# tacacs server host <ip-address> key <plaintextkey>
analytics -1# tacacs server host <ip-address> key 0 <plaintextkey>
analytics -1# tacacs server host <ip-address> key 7 <plaintextkey>
```

Replace **plaintextkey** with a password up to **63** characters in length. This key can be specified either globally or for each host. The first two forms accept a plaintext (literal) key, and the last form accepts a pseudo-encrypted key, such as that displayed with `show running-config`.

It uses the global key value when no key is specified for a given host. An empty key is assumed when no key is specified globally or specified for a given host.

The following example uses the **key 7** option followed by the encrypted string:

```
analytics-1(config-switch)# tacacs server 10.1.1.1 key 7 0832494d1b1c11
```



**Note:** Be careful while configuring TACACS+ to avoid disabling access to the Analytics Node.

## 10.3 Setting up a TACACS+ Server

Refer to your AAA server documentation for further details or instructions on setting up other servers.

After installing the TACACS+ server, complete the following steps to set up authentication and authorization for Analytics Node with the TACACS+ server:

1. Configure users and groups.
2. In the `/etc/tacacs/tac_plus.conf` file, specify the user credentials and group association.

```
# user details
user = user1 {
  member = anet-vsa-admin
  login = des a9qtD2JXeK0Sk
}
```

3. Configure the groups to use one of the AV pairs supported by the Analytics Node (for example, BSN-User-Role=admin for admin users).

```
# group details#
ANET admin group
group = anet-vsa-admin {
  service = exec {
    BSN-User-Role="admin"
  }
}
# BSN read-only group
group = anet-vsa-read-only {
  service = exec {
    BSN-User-Role="read-only"
  }
}
```

4. Configure the TACACS+ server and AAA on the Analytics Node.

```
tacacs server host <IP address> key server's secret>
aaa authentication login default group tacacs+ local
aaa authorization exec default group tacacs+ local
aaa accounting exec default start-stop locals group tacacs+
```

This configuration sets authentication and authorization to first connect to the TACACS+ server to verify user credentials and privileges. It checks the user account locally only when the remote server is unreachable. In this example, accounting stores audit logs locally and sends them to the remote server.

### 10.3.1 Using the Same Credentials for the Analytics Node and Other Devices

The **BSN-User-Role** attribute must be specified as **Optional** in the `tac_plus.conf` file to use the same user credentials to access the Analytics Node and other devices, as shown in the following example.

```
group = group-admin {
  default service = permit
  service = exec {
    optional BSN-User-Role = "admin"
  }
}
```

```
}

```

### 10.3.2 RBAC-based Configuration for Non-default Group User

To create an RBAC configuration for a user in a non-default group, complete the following steps:

1. Create a group **AD1**.

```
group AD1

```

Do not associate with any local users.

2. Use the same group name on the TACACS+ server and associate a user to this group.



**Note:** The attribute should be **BSN-User-Role**, and the value should be the group name.

The following is an example from the open TACACS+ server configuration.

```
group = AD1 {
  service = exec {
    BSN-User-Role="AD1"
  }
}
```

3. After you create the group, associate a user to the group.

```
user = user3 {
  member = AD1
  login = cleartext user3
}
```

4. Click save.

## 10.4 Using RADIUS for Managing Access to the Arista Analytics Node



**Note:** RADIUS does not separate authentication and authorization, so be careful when authorizing a user account with a remote RADIUS server to use the password configured for the user on the remote server.

By default, authentication and authorization functions are set to local while the accounting function is disabled. The only supported privilege levels are as follows:

- **admin:** Administrator access, including all CLI modes and debug options.
- **read-only:** Login access, including most show commands.



**Note:** The **admin** and **recovery** user accounts cannot be authenticated remotely using TACACS. These accounts are always authenticated locally to prevent administrative access from being lost in case a remote AAA server is unavailable.

The **admin** group provides complete access to all network resources, while the **read-only** group provides read-only access to all network resources.

DANZ Monitoring Fabric also supports communication with a remote AAA server (TACACS+ or RADIUS). The following summarizes the options available for each function:

- **Accounting:** local, local and remote, or remote.
- **Authentication:** local, local then remote, remote then local, or remote.
- **Authorization:** local, local then remote, remote then local, or remote.



**Note:** Fallback to local authentication occurs only when the remote server is unavailable, not when authentication fails.

Privileges at the remote TACACS+ server must be configured using the attribute-value pairs shown in the following table:

Supported attribute names	Supported attribute values
BSN-User-Role	admin read-only bigtap-admin bigtap-read-only

The **BSN-AV-Pair** attribute sends CLI command activity accounting to the RADIUS server.

### 10.4.1 Adding a RADIUS Server

Use the following command to specify the remote RADIUS server:

```
radius server host <server-address> [timeout {<timeout>}] [key {{<plaintext>} |  
0 {<plaintext>} | 7 {<secret>}}]
```

For example, the following command identifies the RADIUS server at the IP address **192.168.17.101**:

```
analytics-1(config)# radius server host 192.168.17.101 key admin
```

You can enter this command up to five times to specify multiple RADIUS servers. The Analytics Node tries to connect to each server in the order they are configured.

### 10.4.2 Setting up a FreeRADIUS Server

After installing the FreeRADIUS server, complete the following steps to set up authentication and authorization for the Analytics Node with the RADIUS server:

1. Create the BSN dictionary and add it to the list of used dictionaries.

```
create dictionary /usr/share/freeradius/dictionary.bigswitch with the  
contents below:  
VENDOR      Big-Switch-Networks 37538  
BEGIN-VENDOR Big-Switch-Networks  
ATTRIBUTE   BSN-User-Role 1          string  
ATTRIBUTE   BSN-AVPair    2          string  
END-VENDOR  Big-Switch-Networks
```

2. Include the **bigswitch** dictionary in the RADIUS dictionary file: **/usr/share/freeradius/dictionary**

```
$INCLUDE      dictionary.bigswitch
```

3. Configure a sample user with admin and read-only privileges.

The following is an example that defines and configures a user, opens the user file **/etc/freeradius/users**, and inserts the following entries:

```
"user1"      Cleartext-Password := "passwd"
```

```
BSN-User-Role := "read-only",
```



**Note:** It shows the VSA's association with the user and its privileges. In an actual deployment, a database, and an encrypted password are necessary.

The following example authorizes the **user2** for RBAC group **AD1**:

```
"user2"      Cleartext-Password := "passwd"
             BSN-User-Role := "AD1",
```

#### 4. Configure the RADIUS server and AAA on the Analytics Node.

```
radius server host <IP address> key server's secret>
aaa authentication login default group radius local
aaa authorization exec default group radius local
aaa accounting exec default start-stop group radius local
```

This configuration sets authentication and authorization to first connect to the RADIUS server to verify user credentials and privileges. AAA fallback to local occurs only when the remote server is unreachable. In this example, accounting stores audit logs locally and sends them to the remote server.

#### 5. Add the Analytics Node subnet to the allowed subnets ('clients.conf') on the RADIUS server.

It is required when access to the RADIUS server is limited to allowed clients or subnets. The following is an example of the **clients.conf** file:

```
client anet {
    ipaddr = 10.0.0.0/8
    secret = <server's secret>
}
```

#### 6. Restart the FreeRADIUS service on the server to enable the configuration.

The following is an example accounting record sent from the Analytics Node to the RADIUS server after adding the **BSN-AVPair** attribute to the **/usr/share/freeradius/dictionary.bigswitch** file.

```
s
```

## Creating Watcher Alerts for Machine Learning jobs

The following appendix describes the procedure for creating Watcher alerts for machine learning jobs, emails, and remote Syslog servers.

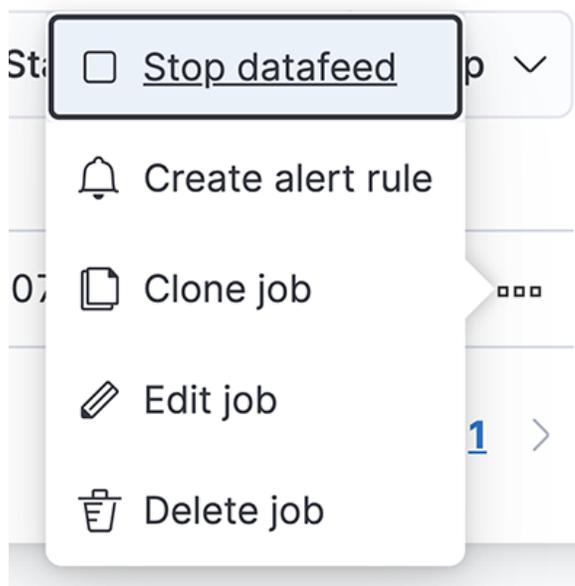
### A.1 Watcher Alert Workaround

**DMF 8.1** uses **Elasticsearch 7.2.0**, where the inter-container functional calls are HTTP-based. However, **DMF 8.3** uses **Elasticsearch version 7.13.0**, which now requires HTTPS-based calls. It would require an extensive change in the system calls used by the Analytics Node (AN), and engineering is working on this effort. Arista recommends the following workaround until the earlier fixes are released.

#### Workaround Summary:

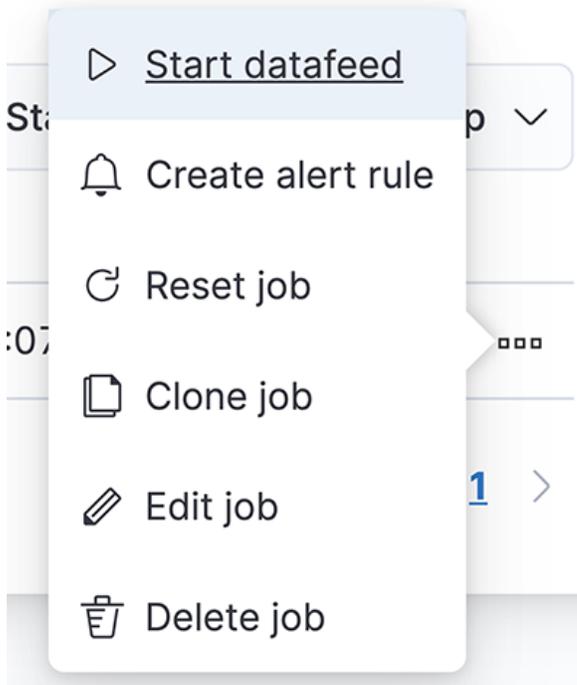
- Create a Watcher manually using the provided template.
  - Configure the Watcher to select the job ID for the ML job that needs to send alerts.
  - Use 'webhook' as the alerting mechanism within the Watcher to send alerts to 3rd party tools like 'Slack.'
1. Access the AN's ML job page and click **Manage Jobs** to list the ML jobs.
  2. If the data feed column shows as **stopped**, skip to **Step 3**. If it says **started**, click the **3 dots** for a particular ML job and **Stop** the data feed for the current ML job.

**Figure A-1: Stop Data Feed**



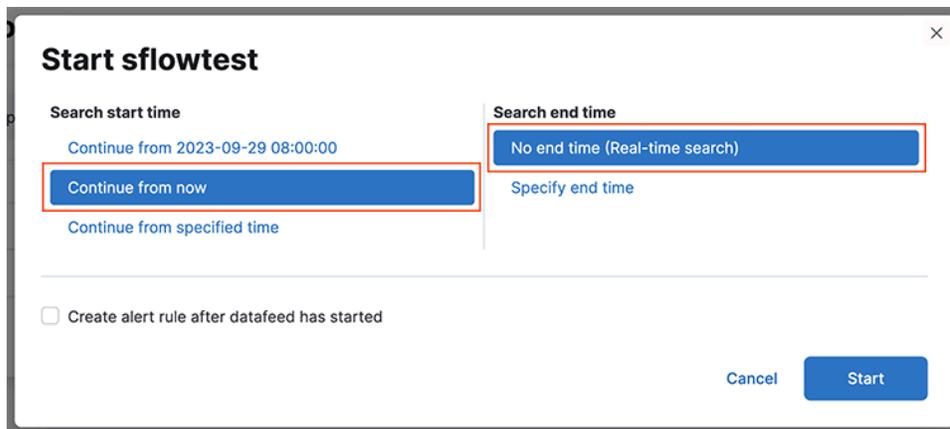
3. After the data feed has stopped, click the **3 dots** and start the data feed.

**Figure A-2: Start Data Feed**



4. Select the options as shown in the diagram below.

**Figure A-3: Job Time Options**



5. Confirm that the data feed has started. Note down the job ID of this ML job.

**Figure A-4: ML Job Characteristics**

The screenshot displays the Elastic ML interface for managing anomaly detection jobs. At the top, there are navigation tabs for Overview, Anomaly Detection, Data Frame Analytics, Data Visualizer, and Settings. The main heading is "Anomaly detection jobs" with a "Refresh" button and a "30 second" interval. Below this, a summary bar shows "Active ML nodes: 1 Total jobs: 1 Open jobs: 1 Closed jobs: 0 Active datafeeds: 1" and a "Create job" button. A search bar contains "id:sflowtest" and a filter menu is set to "Opened".

ID	Description	Processed records	Memo...	Job state	Datafeed state	Latest timestamp
sflowtest		4 ok		opened	started	2023-09-28 16:35:07

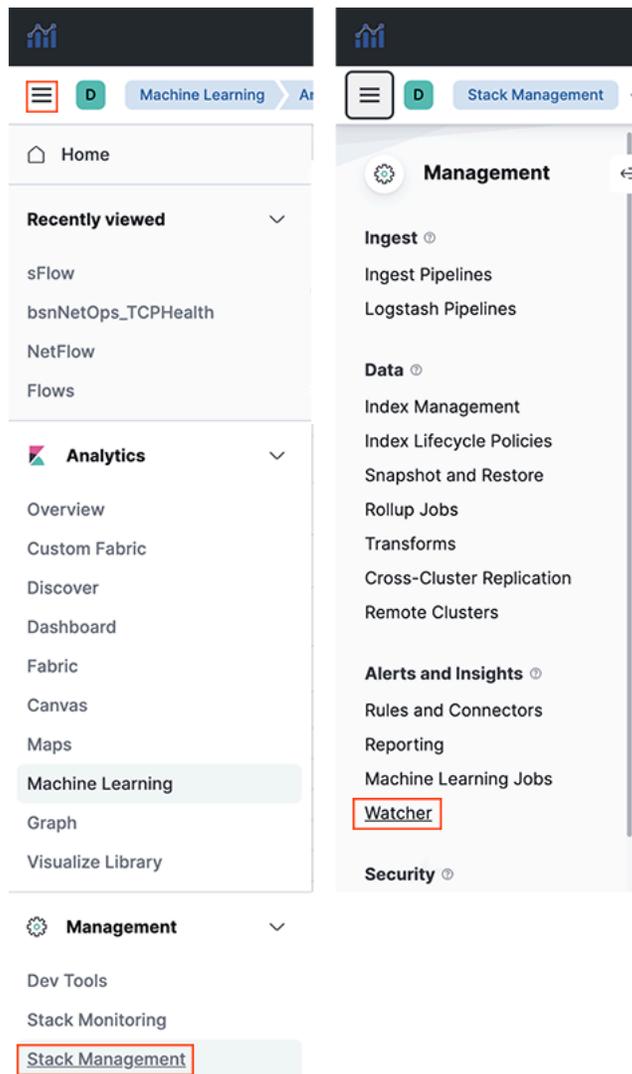
Below the table, the "Job settings" section is active, showing "General" and "Custom settings" tabs. The "General" tab lists various job parameters, with the "job\_id" field highlighted in red and containing the value "sflowtest".

Field	Value
job_id	sflowtest
job_type	anomaly_detector
job_version	7.17.8
create_time	2023-09-29 07:38:47
model_snapshot_id	1695998331
description	
model_snapshot_retention_days	10
daily_model_snapshot_retention_after_days	1
results_index_name	shared
allow_lazy_open	
state	opened
assignment_explanation	
open_time	0s

The "Custom settings" section shows "created\_by" as "single-metric-wizard". The "Node" section shows "name" as "cb3a62ef5e62".

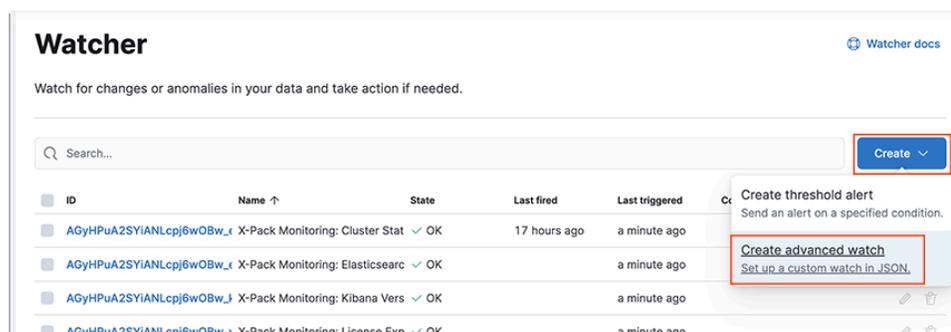
6. Access the **Watchers** page.

**Figure A-5: Access Watchers**



7. Create an advanced Watcher.

**Figure A-6: Create Advanced Watcher**



8. Configure the name of the Watcher (can include whitespace characters), e.g., **Latency ML**.
9. Configure the ID of the Watcher (can be alphanumeric, but without whitespace characters), e.g., **ml\_latency**.

10. Delete the code from the **Watch JSON** section.
11. Copy and paste the following code into the Watcher. Replace the highlighted text according to your environment and your ML job parameters.

```

{
  "trigger": {
    "schedule": {
      "interval": "107s"
    }
  },
  "input": {
    "search": {
      "request": {
        "search_type": "query_then_fetch",
        "indices": [
          ".ml-anomalies-*"
        ],
        "rest_total_hits_as_int": true,
        "body": {
          "size": 0,
          "query": {
            "bool": {
              "filter": [
                {
                  "term": {
                    "job_id": "<use the id of the ML job retrieved in step 6.>"
                  }
                },
                {
                  "range": {
                    "timestamp": {
                      "gte": "now-30m"
                    }
                  }
                }
              ],
              "terms": {
                "result_type": [
                  "bucket",
                  "record",
                  "influencer"
                ]
              }
            }
          }
        }
      },
      "aggs": {
        "bucket_results": {
          "filter": {
            "range": {
              "anomaly_score": {
                "gte": 75
              }
            }
          },
          "aggs": {
            "top_bucket_hits": {
              "top_hits": {
                "sort": [
                  {
                    "anomaly_score": {
                      "order": "desc"
                    }
                  }
                ]
              }
            },
            "source": {
              "includes": [
                "job_id",
                "result_type",
                "timestamp",
                "anomaly_score",
                "is_interim"
              ]
            }
          },
          "size": 1,
          "script_fields": {
            "start": {

```

```

        "script": {
          "lang": "painless",
          "source": "LocalDateTime.ofEpochSecond((doc[\"timestamp\"].value.getMillis()-((doc[\"bucket_span\"].value * 1000)\n * params.padding)) / 1000, 0,ZoneOffset.UTC).toString()+\":00.000Z\"",
          "params": {
            "padding": 10
          }
        },
        "end": {
          "script": {
            "lang": "painless",
            "source": "LocalDateTime.ofEpochSecond((doc[\"timestamp\"].value.getMillis()+((doc[\"bucket_span\"].value * 1000)\n * params.padding)) / 1000, 0,ZoneOffset.UTC).toString()+\":00.000Z\"",
            "params": {
              "padding": 10
            }
          }
        },
        "timestamp_epoch": {
          "script": {
            "lang": "painless",
            "source": "\"doc[\"timestamp\"].value.getMillis()/1000\""
          }
        },
        "timestamp_iso8601": {
          "script": {
            "lang": "painless",
            "source": "\"doc[\"timestamp\"].value\""
          }
        },
        "score": {
          "script": {
            "lang": "painless",
            "source": "\"Math.round(doc[\"anomaly_score\"].value)\""
          }
        }
      }
    }
  },
  "influencer_results": {
    "filter": {
      "range": {
        "influencer_score": {
          "gte": 3
        }
      }
    },
    "aggs": {
      "top_influencer_hits": {
        "top_hits": {
          "sort": [
            {
              "influencer_score": {
                "order": "desc"
              }
            }
          ]
        }
      },
      "_source": {
        "includes": [
          "result_type",
          "timestamp",
          "influencer_field_name",
          "influencer_field_value",
          "influencer_score",
          "isInterim"
        ]
      },
      "size": 3,
      "script_fields": {
        "score": {
          "script": {
            "lang": "painless",
            "source": "\"Math.round(doc[\"influencer_score\"].value)\""
          }
        }
      }
    }
  }
}

```

```

    }
  }
},
"record_results": {
  "filter": {
    "range": {
      "record_score": {
        "gte": 75
      }
    }
  },
  "aggs": {
    "top_record_hits": {
      "top_hits": {
        "sort": [
          {
            "record_score": {
              "order": "desc"
            }
          }
        ]
      },
      "_source": {
        "includes": [
          "result_type",
          "timestamp",
          "record_score",
          "is_interim",
          "function",
          "field name",
          "by_field_value",
          "over_field_value",
          "partition_field_value"
        ]
      },
      "size": 3,
      "script_fields": {
        "score": {
          "script": {
            "lang": "painless",
            "source": ""Math.round(doc["record_score"].value)""
          }
        }
      }
    }
  }
},
"condition": {
  "compare": {
    "ctx.payload.aggregations.bucket_results.doc_count": {
      "gt": 0
    }
  }
},
"actions": {
  "log": {
    "logging": {
      "level": "info",
      "text": "Alert for job [{{ctx.payload.aggregations.bucket_results.top_bucket_hits.hits.hits.0.source.job_id}}] at [{{ctx.payload.aggregations.bucket_results.top_bucket_hits.hits.hits.0.fields.timestamp_iso8601.0}}] score [{{ctx.payload.aggregations.bucket_results.top_bucket_hits.hits.hits.0.fields.score.0}}]"
    }
  },
  "my_webhook": {
    "webhook": {
      "scheme": "https",
      "host": "hooks.slack.com",
      "port": 443,
      "method": "post",
      "path": "<path for slack>",
      "params": {},
      "headers": {

```

```

    "Content-Type": "application/json"
  },
  "body": ""{"channel": "#<slack channel name>", "username": "webhookbot", "text": "Alert
for job [{{ctx.payload.aggregations.bucket_results.top_bucket_hits.hits.hits.0._source
.job_id}}] at [{{ctx.payload.aggregations.bucket_results.top_bucket_hits.hits.hits.0.fields.
timestamp_iso8601.0}}] score [{{ctx.payload.aggregations.bucket_results.top_bucket_hits.hits.h
its.0.fields.score.0}}]", "icon_emoji": ":exclamation:"}""
}
}
}
}
}

```

12. Click **Create Watch** to create the Watcher.

## A.2 Email Alerts and Remote Syslog Server

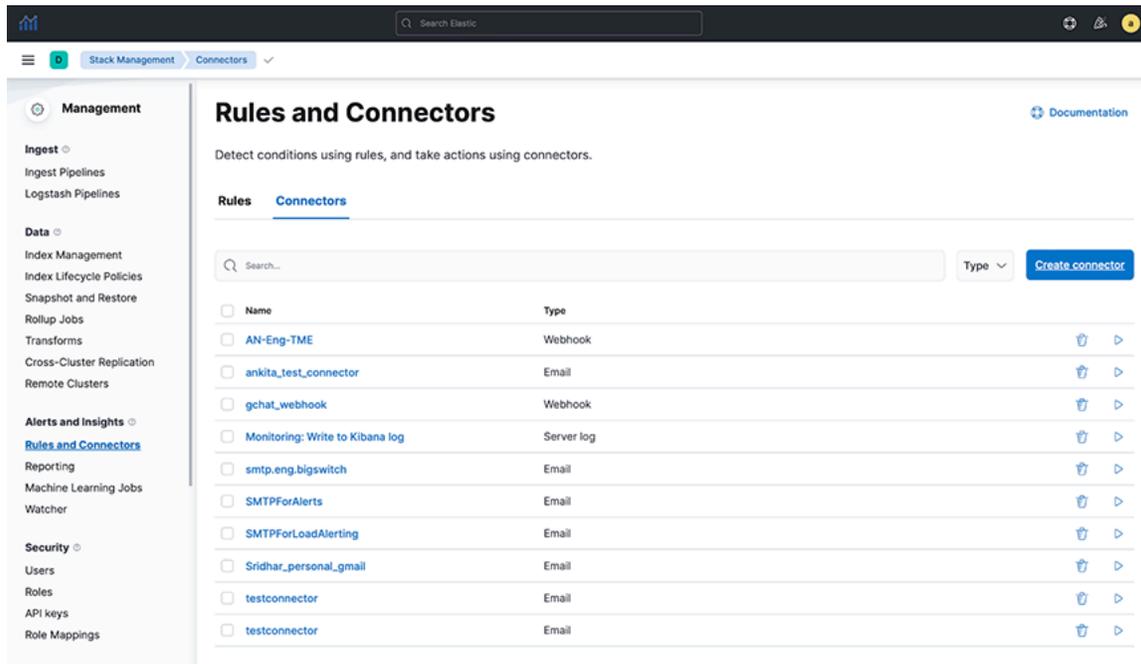
Sending Watcher alerts to email required editing configuration files in the command line and restarting the Elasticsearch container previously.

An update to the Watcher alerts feature creates a simpler configuration method using the Analytics Node UI and supports sending Watcher alerts to remote Syslog servers

### Configuring a Kibana Email Connector

Select an existing Kibana email connector to send email alerts or create a connector by navigating to **Stack Management > Rules and Connectors > Connectors > Create Connectors**. Complete the following steps:

Figure A-7: Rules and Connectors



1. Configure the fields in the **Configuration** tab.

2. Verify the connector works in the **Test** tab.

**Figure A-8: Editing Connector**

The screenshot shows the 'Edit connector' dialog box with the 'Configuration' tab selected. The 'Connector name' field contains 'test\_connector'. Under 'Connector settings', the 'Sender' field contains 'an@arista.com'. The 'Service' dropdown is set to 'Other'. The 'Host' field contains 'smtp.example.com' and the 'Port' field contains '25'. A 'Secure' toggle is currently turned off. Under 'Authentication', the 'Require authentication for this server' toggle is also turned off. At the bottom, there are three buttons: 'Cancel', 'Save', and 'Save & close'.

**Figure A-9: Editing Connector to create action**

The screenshot shows the 'Edit connector' dialog box with the 'Test' tab selected. A vertical progress indicator on the left shows step 1, 'Create an action', which is currently active. The 'To' field has a text input box and 'Cc' and 'Bcc' links. The 'Subject' field has a text input box. The 'Message' field has a large text area. Step 2, 'Run the test', is visible at the bottom but not yet active.

## Configuring a Watch

Configure a Watch using the **Create threshold alert** or **Create advanced watch** option, described in the following instructions.

**Figure A-10: Watcher**

The screenshot shows the Watcher interface with a table of watches. A dropdown menu is open over the 'Create' button, showing two options: 'Create threshold alert' and 'Create advanced watch'.

ID	Name	State	Last fired	Last triggered	Comment
03c44080-837e-442d-b62d-88f82b360c1f	aaa	OK		a few seconds ago	
7a9c22e4-647f-43cc-ac83-dbd9a977110	randomtest	OK		a few seconds ago	
test-watcher	test-watcher	Error	42 minutes ago	42 minutes ago	Execution failing
1547a3c6-c240-4f20-be0f-6c8a831e142d	testwatcher	OK		a few seconds ago	
mialerter		Disabled	3 years ago	3 years ago	
d704ebcd-3afe-4236-82df-956f5a9f4401		OK			
bw_threshold_interfaces		OK	3 years ago	12 minutes ago	
ml-test		OK	22 minutes ago	a minute ago	
673170fe-e755-4509-98ab-07a6a63bb5f		Error	a minute ago	a minute ago	Execution failing
ml-interfaceanalysis		OK		a minute ago	

## Create Threshold Alert

1. Navigate to **Stack Management > Watcher > Create > Create threshold alert** and configure the alert conditions.

**Figure A-11: Creating threshold alert**

The screenshot shows the 'Create threshold alert' configuration page. It includes a name field, indices to query, time field, and run watch every settings. Below these is a section for matching conditions, which currently shows 'No data'.

**Create threshold alert**

Send an alert when your specified condition is met. Your watch will run every 1 minute.

Name: test watcher

Indices to query: .kibana\*

Time field: alert.createdAt

Run watch every: 1 minute

**Match the following condition**

```
WHEN count() OVER all documents IS ABOVE 1000 FOR THE LAST 5 minutes
```

No data  
Your index and condition did not return any data.

2. Add a webhook action with the following fields.
  - **Method:** POST
  - **Scheme:** HTTP
  - **Host:** 169.254.16.1
  - **Port:** 8000
  - Specify the **Body** field as follows:
    - **Sending Watcher alerts by email:** Enter the required fields: **to**, **subject**, **message**, and **kibana\_email\_connector**. Multiple entries in the **to** field require a comma-separated list of email

addresses wrapped in quotes. The `kibana_email_connector` field references an existing Kibana email connector.

- **Sending Watcher alerts to a remote Syslog server:** Enter the required fields: **message**, **protocol**, **primary\_syslog\_ip**, and **primary\_syslog\_port**. If a second Syslog server should receive alerts, include **backup\_syslog\_ip** and **backup\_syslog\_port**.

**Figure A-12: Performing Action for Webhook**

Perform 1 action when condition is met Add action ▾

Webhook ✕

Method: POST ▾ Scheme: http ▾ Host: 169.254.16.1 Port: 8000 ▾ Path (optional): /

Username (optional): Password (optional):

Body

```
{
  "message": "Watch [{{ctx.metadata.name}}] has exceeded the threshold",
  "subject": "subject",
  "to": "recipient@domain.com",
  "kibana_email_connector": "SMTPForAlerts",
  "protocol": "UDP",
  "primary_syslog_ip": "10.240.145.3",
  "primary_syslog_port": 514,
  "backup_syslog_ip": "10.240.145.4",
  "backup_syslog_port": 514
}
```

Send request

✓ Create alert Cancel Show request

- The **Path**, **Username**, and **Password** fields do not need to be specified.
3. Test the webhook action using **Send Request** before selecting **Create alert**. Depending on the configuration:
- Verify the receipt of an email at the configured recipient address.
  - Verify the receipt of a syslog message on the remote Syslog server.

## Create Advanced Watch

1. Navigate to **Stack Management > Watcher > Create > Create advanced watch** and fill out the Name and ID of the Watch.

Figure A-13: Editing Advanced Watch

**Create advanced watch**

[Edit](#) [Simulate](#)

Name (optional)

ID

Watch JSON (API syntax) [?](#)

```

- {
-   "trigger": {
-     "schedule": {
-       "interval": "30m"
-     }
-   },
-   "input": {
-     "search": {
-       "request": {
-         "body": {
-           "size": 0,
-           "query": {
-             "match_all": {}
-           }
-         },
-         "indices": [
-           "*"
-         ]
-       }
-     },
-     "condition": {
-       "compare": {
-         "ctx.payload.hits.total": {
-           "gte": 10
-         }
-       }
-     },
-     "actions": {
-       "my-logging-action": {
-         "logging": {
  
```

[Show request](#)

2. For the Watch JSON field, the following JSON template configures the forwarding of alerts to email and remote Syslog servers. Configure the alert condition under the **input** and **condition** fields. Replace these values with any custom alert condition using the Elastic Painless scripting language. The configuration for forwarding alerts to email and remote Syslog servers is under the **actions** field.

```

{
  "trigger": {
    "schedule": {
      "interval": "1m"
    }
  },
  "input": {
    "http": {
      "request": {
        "scheme": "https",
        "host": "<host>",
        "port": 443,
        "method": "get",
        "path": "/_cluster/health",
        "params": {},
        "headers": {
          "Content-Type": "application/json"
        }
      },
      "auth": {
        "basic": {
          "username": "<user>",
          "password": "<password>"
        }
      }
    }
  }
}
  
```

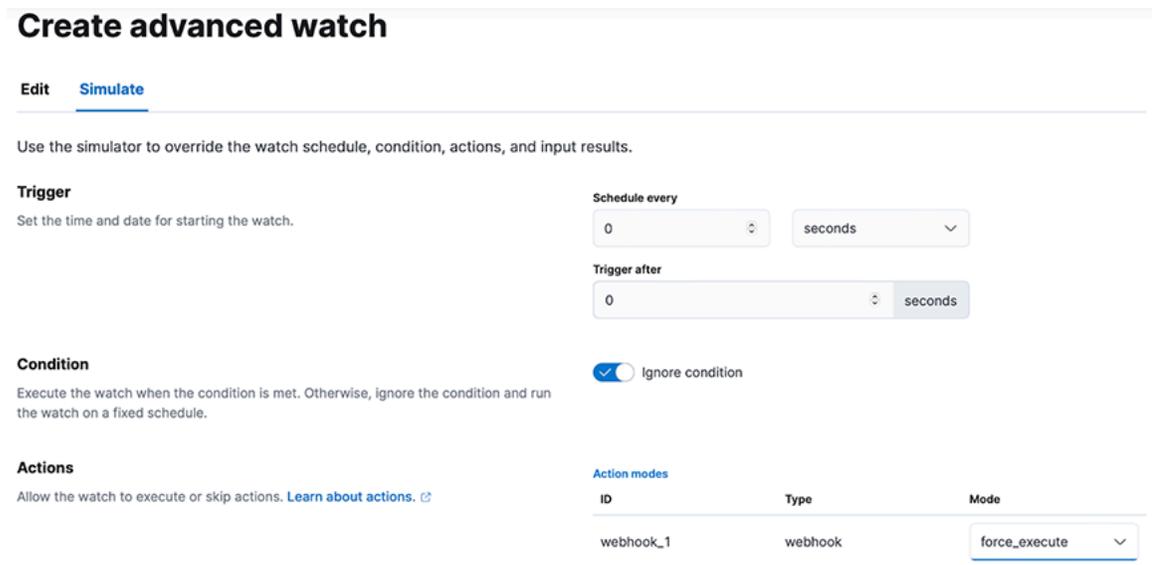
```

    }
  },
  "condition": {
    "script": {
      "source": "ctx.payload.status == 'green'",
      "lang": "painless"
    }
  },
  "actions": {
    "webhook_1": {
      "webhook": {
        "host": "169.254.16.1",
        "port": 8000,
        "method": "post",
        "scheme": "http",
        "body": "{\n\"message\": \"The Elasticsearch cluster status is
        {{ctx.payload.status}}\", \n\"kibana_email_connector\": \"<existing-email-
        connector>\", \n\"to\": \"recipient@domain.com\", \n\"subject\": \"Elasticsear
        ch cluster status alert\", \n\"protocol\": \"UDP\", \n\"primary_syslog_ip\":
        \"<remote-syslog-ip>\", \n\"primary_syslog_port\": <remote-syslog-port>,
        \n\"backup_syslog_ip\": \"<remote-syslog-ip>\", \n\"backup_syslog_port\":
        <remote-syslog-port>}"
      }
    }
  }
}

```

3. (Optional) To simulate the Watch, you can configure the fields in the Simulate Tab. The webhook action mode must be set to **force\_execute**.

**Figure A-14: Simulating Advanced Watch**



## Troubleshooting

- If the email alert fails, verify that the value of the **kibana\_email\_connector** field matches the name of a Kibana email connector and that this email connector works in the **Test** tab.

## Limitations

- Remote Syslog messages require UDP. TCP is not supported currently.

## A.3 Enabling Secure Email Alerts through SMTP Setting

Refresh the page to view the updated **SMTP Settings** fields.

The following is an example of the UI SMTP Settings in previous releases:

**Figure A-15: SMTP Setting**

The screenshot shows the 'Configure Alerts' interface with the 'SMTP Settings' section. The settings are as follows:

Field	Value
Server Name	
User	
Password	
Recipients	recipient-a, recipient-b, ...
Sender	
Timezone	eg. America/Los_Angeles, UTC, ..
Dedupe Interval (m)	eg. 600

Buttons: Apply & Test, Cancel

After upgrading the Analytics Node from an earlier version to the **DMF 8.6.\*** version, the following changes apply:

- Server **Name**, **User**, **Password**, **Sender**, and **Timezone** no longer appear in the SMTP Settings.
- A new field, **Kibana Email Connector Name**, has been added to SMTP Settings.
- The system retains **Recipients** and **Dedupe Interval** and their respective values in SMTP Settings.
- If previously configured SMTP settings exist:
  - The system automatically creates a Kibana email connector named **SMTPForAlerts** using the values previously specified in the fields Server Name, User (optional), Password (optional), and Sender.
  - The **Kibana Email Connector Name** field automatically becomes **SMTPForAlerts**.

The following settings appear in the UI after the upgrade to the **DMF 8.6.\*** version:

**Figure A-16: Upgraded SMTP Setting**

The screenshot shows the 'Configure Alerts' interface with the 'SMTP Settings' section after the upgrade. The settings are as follows:

Field	Value
Recipients	recipient@domain.com
Dedupe Interval (m)	30
Kibana Email Connector Name	SMTPForAlerts

Buttons: Apply & Test, Cancel

## Troubleshooting

When **Apply & Test**, do not send an email to the designated recipients, verify the recipient email addresses are comma-separated and spelled correctly. If it still doesn't work, verify the designated Kibana email connector matches the name of an existing Kibana email connector. Test that connector by navigating to **Stack Management > Rules and Connectors > Connectors**, selecting the connector's name, and sending a test email in the **Test** tab.

## References

---

### B.1 Related Documents

The following documentation is available for *Arista Analytics 8.6.0*:

- *Arista Analytics User Guide*
- *Arista Analytics Deployment Guide*