

Juniper Apstra Flow Installation and Upgrade Guide

Published
2024-09-17

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Juniper Apstra Flow Installation and Upgrade Guide
Copyright © 2024 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide | v

1

Juniper Apstra Flow Overview

Apstra Flow Overview | 2

2

Get Started with Juniper Apstra Flow

Supported Hypervisors and Versions | 4

Apstra Flow Scaling Considerations | 4

Supported Protocols and Devices | 5

3

Install and Configure Apstra Flow

Configure Your Network and Flow Collector | 7

Configure Your Network | 7

Configure the Flow Collector | 9

Create an OpenSearch Multinode Cluster (Optional) | 10

Create a Multinode Cluster | 11

Multinode Roles and Configurations | 13

Apply Your License | 13

Apply Your License in Apstra | 14

Restart the Flow Services | 15

Import the Junos Apstra Flow Configlet | 15

Import the Configlet | 16

Example: Junos Apstra Flow Configlet | 17

Configure SNMP for Interface Name Enrichment | 18

4

Access the Apstra Flow Dashboards

Launch the Apstra Flow Dashboard | 21

Access the Apstra Flow Dashboard from the Apstra GUI | 25

5

Upgrade Apstra Flow

Upgrade Apstra Flow on the Same VM (In-Place) | 28

| Upgrade Asptra Flow | 28

About This Guide

This guide describes how to install and configure Juniper Apstra Flow on your network.

We'll walk you through how to configure your network and flow collector, and how to apply your license in Apstra. Then, we'll show you how to launch the Apstra Flow dashboards to analyze and visualize your data.

1

CHAPTER

Juniper Apstra Flow Overview

[Apstra Flow Overview](#) | 2

Apstra Flow Overview

Juniper's Apstra Flow feature is a robust solution for collecting and analyzing data center network flow traffic. This feature streamlines the gathering of network traffic flows and telemetry by offering a seamless integration with your organization-specific information. By doing so, it delivers unmatched visibility and insight into your network traffic.

The Apstra Flow dashboards, with enriched data and analytics, provide you with a holistic understanding of the network for various critical use cases, including:

- **Performance and availability**

Provides granular information about network traffic flows, congestion, high latency, and packet loss.

- **Capacity planning and cost control**

Enables you to implement strategies to optimize the flow of network traffic, ensuring the most efficient use of available resources.

- **Security**

Improves security by detecting and responding to threats more effectively while maintaining compliance with regulatory requirements.

With Apstra Flow, you gain a powerful toolset to optimize and fortify your network infrastructure.

NOTE: Apstra Flow is a feature in the Apstra Premium tier licensing plan. This feature is available only if you are an Apstra premium customer. For more information, see the [Juniper Licensing User Guide](#).

2

CHAPTER

Get Started with Juniper Apstra Flow

Supported Hypervisors and Versions | 4

Apstra Flow Scaling Considerations | 4

Supported Protocols and Devices | 5

Supported Hypervisors and Versions

Apstra Flow Data supports the following hypervisors and versions:

Table 1: Supported Hypervisors and Versions

Hypervisor	Supported Versions
VMware vSphere	7.0, 6.7, 6.5, 6.0
Microsoft Hyper-V	Windows Server 2016 Datacenter Edition
Linux KVM	Ubuntu 22.04, 18.04 LTS

Apstra Flow Scaling Considerations

The Apstra Flow virtual images are available in three sizing options, plus an additional option for creating custom extra-large deployments. [Table 1 on page 4](#) describes the available sizing options and memory requirements.

Table 2: Apstra Flow Scaling Considerations

*Image Size	Required Memory Size	Number of Supported Devices
Small Up to 5,000 FPS	8 vCPU, 32 GB memory, 4 TB disk	Typically, up to 25 devices at an average of 200 FPS per device.
Medium Up to 10,000 FPS	16 vCPU, 64 GB memory, 4 TB disk	Typically, up to 50 devices at an average of 200 FPS per device
Large Up to 15,000 FPS	24 vCPU, 64 GB memory, 4 TB disk	Typically, up to 75 devices at an average of 200 FPS per device

Table 2: Apstra Flow Scaling Considerations (Continued)

*Image Size	Required Memory Size	Number of Supported Devices
X-Large (custom)	N/A	Greater than 15,000 FPS. If you have an extra-large or custom deployment, we recommend that you create a multinode cluster. See " Create an OpenSearch Multinode Cluster (Optional) " on page 10.

*Image size is calculated at an average number of flows per second (FPS) that the collector receives.

Supported Protocols and Devices

Apstra Flow Data supports the following protocols and devices:

- sFlow, NetFlow v1, v5, v6, v7, v9, IPFIX, and IFA protocols
- Juniper Junos OS devices
- Juniper Junos Evolved OS devices

Note: Sending sFlow through the management interface on Junos OS Evolved devices is not supported.

- Arista EOS, Cisco NXOS, and Enterprise SONIC devices

See the appropriate vendor documentation for device-specific details.

3

CHAPTER

Install and Configure Apstra Flow

[Configure Your Network and Flow Collector](#) | 7

[Create an OpenSearch Multinode Cluster \(Optional\)](#) | 10

[Apply Your License](#) | 13

[Import the Junos Apstra Flow Configlet](#) | 15

Configure Your Network and Flow Collector

IN THIS SECTION

- [Configure Your Network | 7](#)
- [Configure the Flow Collector | 9](#)

NOTE: Before you begin the configuration, make sure that you download the image file for your virtualization platform. See your platform documentation for details.

NOTE: You can install KVM with Virtual Machine Manager or with the CLI.

Let's get started!

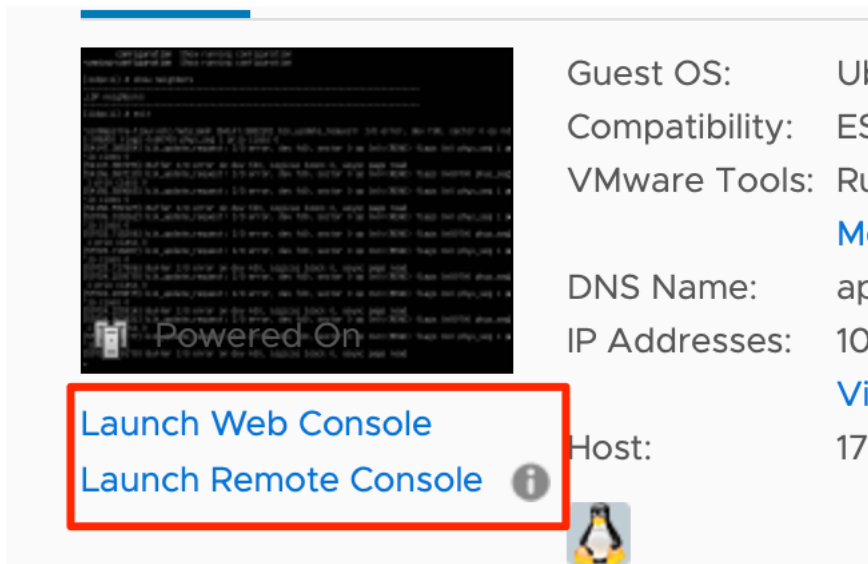
Configure Your Network

NOTE: If you have a DHCP server, by default, the Apstra Flow VM automatically obtains an IP address. To assign a static IP address instead, SSH to the VM and skip to Step 2 in this procedure. Otherwise, start at Step 1.

1. If the console is already started, click the play button



at the top of the VM summary.



2. Login to the CLI. The default credentials for the Apstra Flow console are: user=apstra and password=apstra. You'll want to change this password after the VM is deployed.
3. Change the directory to /etc/netplan.

```
apstra@apstra-flow:~$ cd /etc/netplan/
apstra@apstra-flow:/etc/netplan$ ls
00-installer-config.yaml 01-netcfg.yaml
```

4. Add an IP address to your network interface. Edit the 01-netcfg.yaml file as follows:

```
apstra@apstra-flow:/etc/netplan$ sudo vi 01-netcfg.yaml
```

```
#
network:
  version: 2
  ethernets:
    eth0:
      dhcp4: false
      addresses:
        - 10.19.10.100/24 <replace with static IP, if desired>
        # - 192.0.2.1
        # search: []
      routes:
        - to: 10.11.0.0/16 <add your DNS settings>
          via: 10.19.10.1 <-
```

```

routes:
- to: 0.0.0.0/0
  via: 172.25.90.1 <configure the default route>

```

5. Enter `sudo netplan apply` to power up your network.

```
apstra@apstra-flow:/etc/netplan$ sudo netplan apply
```

```
** (generate:16454): WARNING **: 14:16:45.231: Permissions for /etc/netplan/00-installer-
config.yaml are too open.
```

```
Netplan configuration should NOT be accessible by others.
```

```
** (generate:16454): WARNING **: 14:16:45.231: Permissions for /etc/netplan/01-netcfg.yaml
are too open.
```

```
Netplan configuration should NOT be accessible by others.
```

```
(you can ignore permissions errors, or fix if desired)
```

You can now SSH to the VM with your configured IP address and Apstra credentials, as shown in the following example:

```

username@username-mbp ~ % ssh apstra@172.25.90.2
apstra@172.25.90.2's password:
Welcome to Ubuntu 22.04.3 LTS (GNU/Linux 5.15.0-86-generic x86_64)

```

NOTE: Optionally, you can create a VM with a multinode cluster for scaling and high availability. See "[Create an OpenSearch Multinode Cluster \(Optional\)](#)" on page 10 for more information.

Configure the Flow Collector

You can configure the collector to specify a hostname or IP address, the port the collector is listening on, and username and password for the Apstra controller.

To configure the flow collector:

1. SSH to the Flow VM with the default VM username and password (apstra/apstra).

2. Modify the `/etc/juniper/flowcoll.yml` file as follows:

- Set `EF_JUNIPER_APSTRA_API_HOSTNAME` to the Apstra controller IP or hostname (for example: 10.0.0.3)
- Set `EF_JUNIPER_APSTRA_PORT`: 443 (this is the default port)
- Set `EF_JUNIPER_APSTRA_API_USERNAME` to the Apstra controller username (for example: admin)
- Set `EF_JUNIPER_APSTRA_API_PASSWORD` to the Apstra controller password (for example: admin)

NOTE: By default, Apstra Flow reports the interface index values, like `index:14`. To enrich the reporting with actual interface names, like `xe-0/0/1`, configure Apstra Flow to query devices through SNMP using the following settings:

- `EF_PROCESSOR_ENRICH_NETIF_SNMP_COMMUNITIES`: public
- `EF_PROCESSOR_ENRICH_NETIF_SNMP_ENABLE`: "true"

You completed configuring your network and flow collector. Next, continue to ["Apply Your License" on page 13](#).

Create an OpenSearch Multinode Cluster (Optional)

IN THIS SECTION

- [Create a Multinode Cluster | 11](#)
- [Multinode Roles and Configurations | 13](#)

If you have an extra-large or custom deployment that requires more scale than those referenced in ["Apstra Flow Scaling Considerations" on page 4](#), we recommend that you create an OpenSearch multinode cluster. A multinode cluster is a collection of nodes, that work together as a single unit. A multinode cluster can scale to hundreds of nodes, ensures high availability, and prevents downtime.

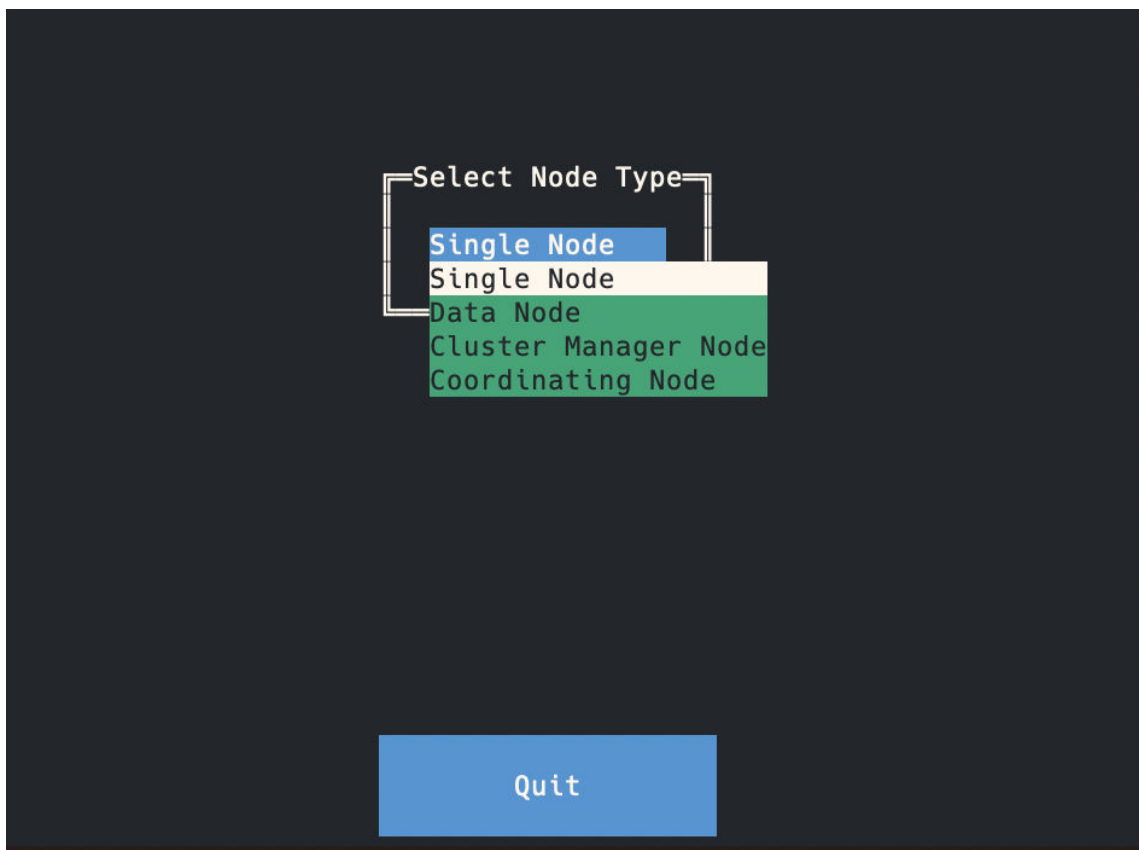
Create a Multinode Cluster

To create a multinode OpenSearch cluster, follow these steps:

1. Configure the node type and associated roles.

- a. From the CLI, run `/usr/local/bin/startup`.
- b. Choose the node type from the TUI (text-based UI).

You can choose from Data Node, Cluster Manager Node, or Coordinating Node. For descriptions about the different node types, see ["Multinode Roles and Configurations" on page 13](#).



- c. Select `Quit`.

When the TUI closes, the node name and node type are automatically set.

- d. Repeat steps 1 through 3 for each node in your configuration.

2. Configure the cluster.

Edit the `opensearch.yaml` file to set the IP address on all nodes in the cluster. For example:

```
# set the node names for the manager nodes
```



```
cluster.initial_cluster_manager_nodes: ["node-1", "node-2", "node-3"]
# Set the ip addresses of the other nodes in the cluster
discovery.seed_hosts: ["node-1-ip", "node-2-ip", "node-3-ip"]
```

- Restart OpenSearch for the changes to take affect.

```
sudo systemctl restart opensearch
```

- Run the following command to verify that cluster was successfully created.

```
curl -X GET "http://node-1-ip:9200/_cluster/health?pretty"
```

In the results output, the status has three possible values as described below. In this example, the green signifies that the cluster is healthy.

- green: All primary and replica shards are active.
- yellow: All primary shards are active, but some replica shards are unassigned.
- red: Some primary shards are not active.

```
{
  "cluster_name": "my-cluster",
  "status": "green",
  "timed_out": false,
  "number_of_nodes": 3,
  "number_of_data_nodes": 3,
  "active_primary_shards": 10,
  "active_shards": 20,
  "relocating_shards": 0,
  "initializing_shards": 0,
  "unassigned_shards": 0,
  "delayed_unassigned_shards": 0,
  "number_of_pending_tasks": 0,
  "number_of_in_flight_fetch": 0,
  "task_max_waiting_in_queue_millis": 0,
  "active_shards_percent_as_number": 100.0
}
```

Multinode Roles and Configurations

- **Data node**

Stores data and participates in the cluster's indexing and search capabilities. Start with 2+ nodes and add more as data volumes grow.

- **Cluster manager node**

Manages the overall operation of a cluster and keeps track of the cluster state. Use odd numbers (3 or 5) for quorum. Avoid using an even number, which can cause instability.

- **Coordinating (client) node**

Manages search and indexing requests, and helps distribute the load across the cluster. Use 2 to 3 nodes to offload data nodes.

Scaling Considerations

- **Extra large or custom sizing options**

We recommend that you transition to a multinode cluster if you are processing over 1 TB of data or greater than 15,000 FPS. You might also want to consider a multinode cluster if you require OpenSearch to be highly available.

- **Odd versus even number of nodes**

For cluster manager nodes, always maintain an odd number to ensure a quorum. You can configure data, ingest and client nodes based on workload without this restriction.

- **Minimum node counts**

For a production environment, a minimum of three nodes is recommended to ensure fault tolerance (three cluster manager nodes and at least two data nodes).

Apply Your License

IN THIS SECTION

- [Apply Your License in Apstra | 14](#)

Apply Your License in Apstra

IN THIS SECTION

- [Restart the Flow Services | 15](#)

By now, you should have received your activation code and activated your license from the [Juniper Agile Licensing Portal](#) (see the [Juniper Licensing User Guide](#) for activation instructions). You now need to apply your license in the Apstra GUI.

To apply your license:

1. Login to the Apstra GUI using the credentials you set when you first configured the Apstra server.
2. From the left navigation menu, select **Platform > Licenses**, then click **Create License**.

ID	Feature Name	Description	License Key	Start Date	Expiration Date	Actions
TRIALJUNOS	Apstra Premium	Premium feature set - more than 3 blueprints, non-Juniper fabrics, policy assurance, flow	[Redacted]	2023-10-15, 19:00:00	2023-12-15, 18:00:00	[Actions]

3. Copy your license key into the **License Key** text box, then click **Create**.

Create License

License Key

Enter your License Key here

Create Another?

4. Verify your license. If the license was successfully applied, the **Feature Name** should indicate **Apstra Premium**.

The screenshot shows the 'Platform Licenses' page in Apstra. At the top right, there is a 'Create License' button. Below it is a search bar and a pagination control showing '1-1 of 1'. The main content is a table with the following columns: ID, Feature Name, Description, License Key, Start Date, Expiration Date, and Actions. One license is listed with ID 'JUNOS161254970', Feature Name 'Apstra Premium', and Description 'Premium feature set - more than 3 blueprints, non-Juniper fabrics, policy assurance, flow'. The 'Feature Name' cell is highlighted with a red border.

ID	Feature Name	Description	License Key	Start Date	Expiration Date	Actions
JUNOS161254970	Apstra Premium	Premium feature set - more than 3 blueprints, non-Juniper fabrics, policy assurance, flow		2023-09-06, 19:00:00	2028-09-05, 19:00:00	

Restart the Flow Services

After you apply your license in Apstra, restart your flow services so that the services can pick up the license.

To restart the services:

1. SSH to the Flow VM. The default login credentials are: apstra/apstra.

2. Restart the collector service.

```
sudo systemctl restart flowcoll
```

3. Restart the OpenSearch service.

```
sudo systemctl restart opensearch
```

4. Restart the OpenSearch dashboard service.

```
sudo systemctl restart opensearch-dashboards
```

5. Verify your license.

Enter the IP address of your VM.

```
http://<flow-vm-ip>:8080/metrics
```

If your license was successfully applied, the `license_units` value will be 16.

Continue to ["Import the Junos Apstra Flow Configlet" on page 15](#).

Import the Junos Apstra Flow Configlet

IN THIS SECTION

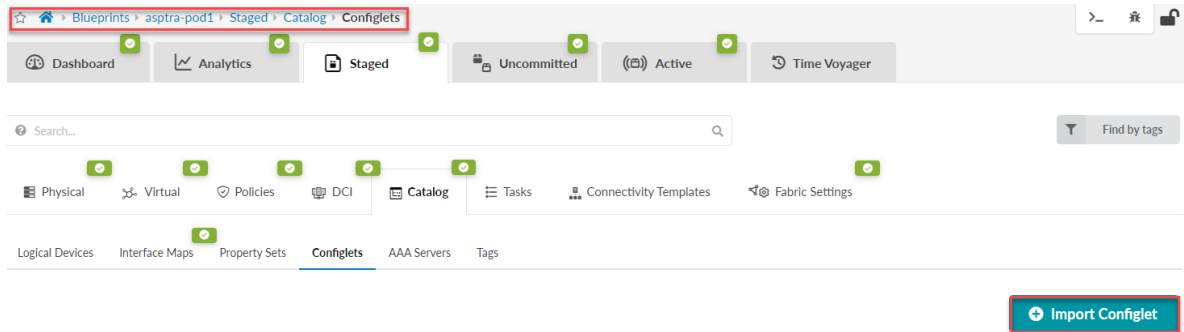
- [Import the Configlet | 16](#)
- [Example: Junos Apstra Flow Configlet | 17](#)
- [Configure SNMP for Interface Name Enrichment | 18](#)

Import the Configlet

A standard Junos sFlow configlet is provided with your product. You'll want to clone this default configlet, modify it to your needs, and then import it into your blueprint from the Apstra GUI (see ["Example: Junos Apstra Flow Configlet" on page 17](#)).

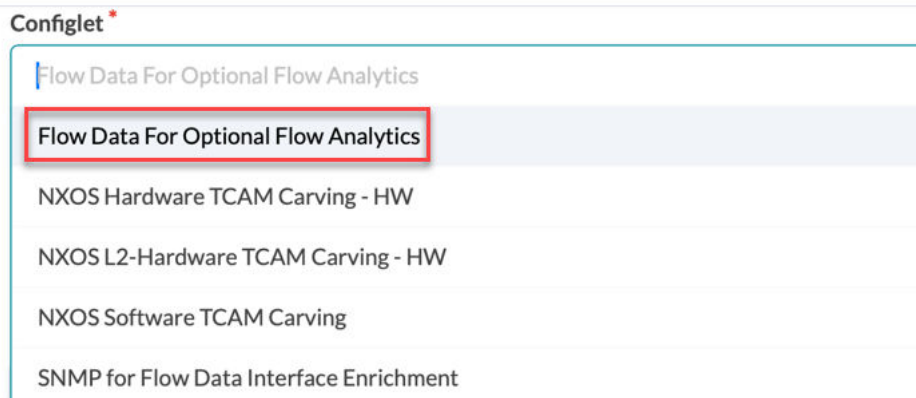
To import the flow data configlet into Apstra:

1. From the left navigation menu in the Apstra GUI, navigate to **Blueprints > Analytics > Staged**. From the **Catalog** tab, select **Configlets**, then click **Import Configlet**.



2. From the **Configlet** drop-down menu, select **Flow Data for Optional Flow Analytics**. Select the devices you want to apply the configlet to, then click **Import Configlet**.

Import Configlet from Global Catalog



3. Edit the property set to add the collector IP address of your Flow Data VM, then click **Update**.

Edit Property Set

Name *

Flow Data For Optional Flow Analytics

Input Type

Editor Builder

Values *

JSON YAML

1 collector_ip: IP_OF_FLOW_VM
2

Update

- Under **Property Set**, select **Flow Data For Optional Flow Analytics**, then click **Import Property Set**.

Import Property Set from Global Catalog

Property Set *

Flow Data For Optional Flow Analytics

Flow Data For Optional Flow Analytics

SNMP for Flow Data Interface Enrichment

- Finally, commit your configuration to import the configlet to your Junos devices.

Good job! You completed the Apstra Flow Data installation. Next, continue to ["Launch the Apstra Flow Dashboard" on page 21](#).

Example: Junos Apstra Flow Configlet

The following example shows the Junos Apstra Flow configlet provided with your product.

This configlet contains a static routing option to route the flow traffic over the management Junos RPF interface. You can also specify the sFlow polling and sampling intervals, collector IP address, UDP port and interfaces you want to collect data on.

```
routing-options {
  static {
    route 10.28.36.6/32 next-table mgmt_junos.inet.0;
  }
}
protocols {
  sflow {
```

```

polling-interval 10;
sample-rate {
    ingress 1024;
    egress 1024;
}
source-ip 10.28.36.12;
collector 10.28.36.6 {
    udp-port 6343;
}
interfaces ge-0/0/0;
interfaces ge-0/0/1;
interfaces ge-0/0/2;
}
}
}

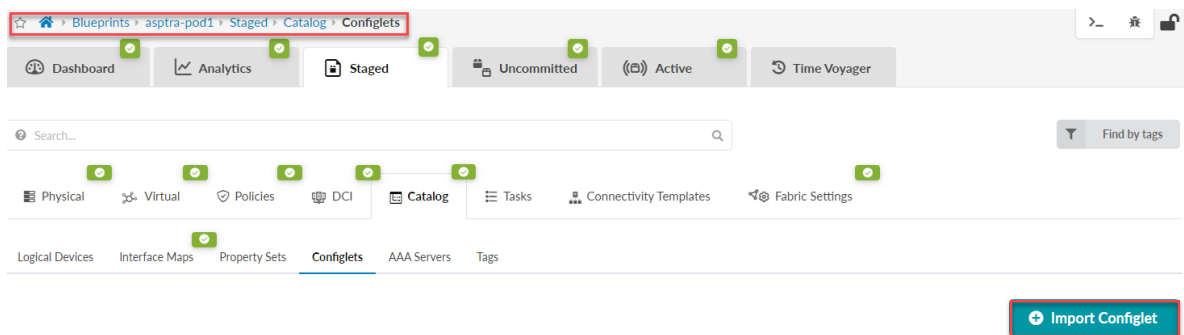
```

Configure SNMP for Interface Name Enrichment

Optionally, you can display the actual interface names (for example, xe-0/0/1) instead of the interface indexes values. To do this, Apstra Flow needs to query the devices through SNMP for this type of enrichment.

Perform the following steps to configure SNMP:

1. From the left navigation menu in the Apstra GUI, navigate to **Blueprints > Analytics > Staged**. From the **Catalog** tab, select **Configlets**, then click **Import Configlet**.



2. From the **Global Catalog**, select **SNMP for Flow Data Interface Enrichment**. Select the devices you want to apply the configlet to, then click **Import Configlet**.

Import Configlet from Global Catalog

Configlet *

- Flow Data For Optional Flow Analytics
- Flow Data For Optional Flow Analytics
- NXOS Hardware TCAM Carving - HW
- NXOS L2-Hardware TCAM Carving - HW
- NXOS Software TCAM Carving
- SNMP for Flow Data Interface Enrichment**

3. Edit the property set to add the community string for your SNMP configuration, then click **Update**.

Edit Property Set

Name *

SNMP for Flow Data Interface Enrichment

Input Type

Editor Builder

Values *

JSON **YAML**

```
1 snmp_community: public
2
```

Update

4. Under **Property Set**, select **SNMP for Flow Data Interface Enrichment**, then click **Import Property Set**.

Import Property Set from Global Catalog

Property Set *

- Flow Data For Optional Flow Analytics
- Flow Data For Optional Flow Analytics
- SNMP for Flow Data Interface Enrichment**

5. Finally, commit your configuration to import the configlet to your Junos devices.

4

CHAPTER

Access the Apstra Flow Dashboards

[Launch the Apstra Flow Dashboard](#) | 21

Launch the Apstra Flow Dashboard

IN THIS SECTION

- [Access the Apstra Flow Dashboard from the Apstra GUI | 25](#)

Congratulations! You successfully installed and configured Apstra Flow on your network. Now you're ready to login to the Apstra Flow dashboard to begin analyzing and visualizing your data.

NOTE: You can also access the Apstra Flow dashboard directly from your blueprint in the Apstra GUI. See "[Access the Apstra Flow Dashboard from the Apstra GUI](#)" on page 25 for instructions.

To launch the Apstra Flow dashboard:

1. From your browser, enter the IP address or hostname of your Apstra Flow VM.

`https://<ip-or-hostname-of-flow-data-vm>:5601/`

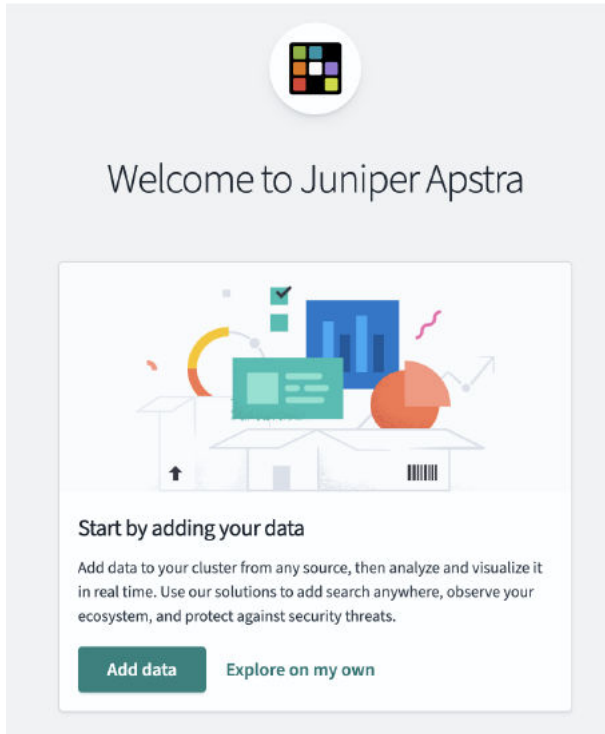
2. Login to the Apstra Flow dashboard.

The default credentials are:

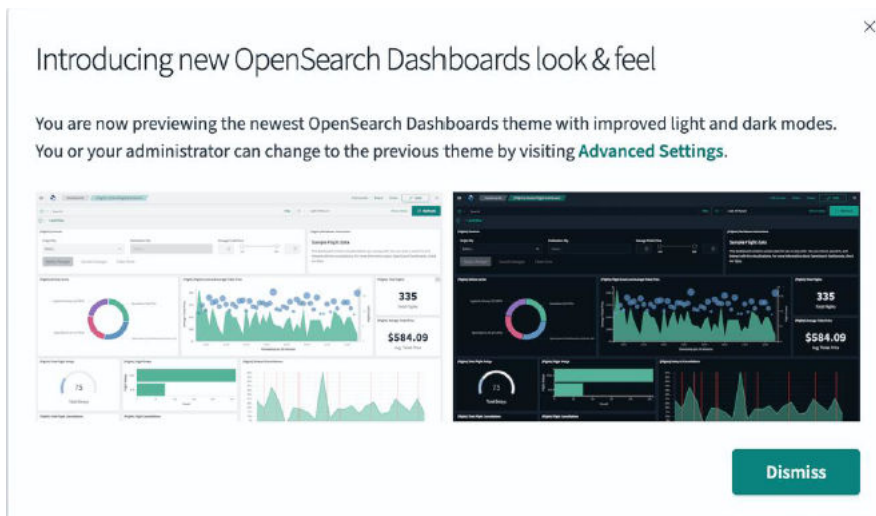
- Username: admin
- Password: Apstra-Flow5

NOTE: We recommend that you change your password after you log in.

3. From the Welcome page, select **Explore on my own**.



4. From the OpenSearch Dashboards page, click **Dismiss**.



5. Select **Global** as your tenant, then click **Confirm**.

Select your tenant

Tenants are useful for safely sharing your work with other OpenSearch Dashboards users. You can switch your tenant anytime by clicking the user avatar on top right.

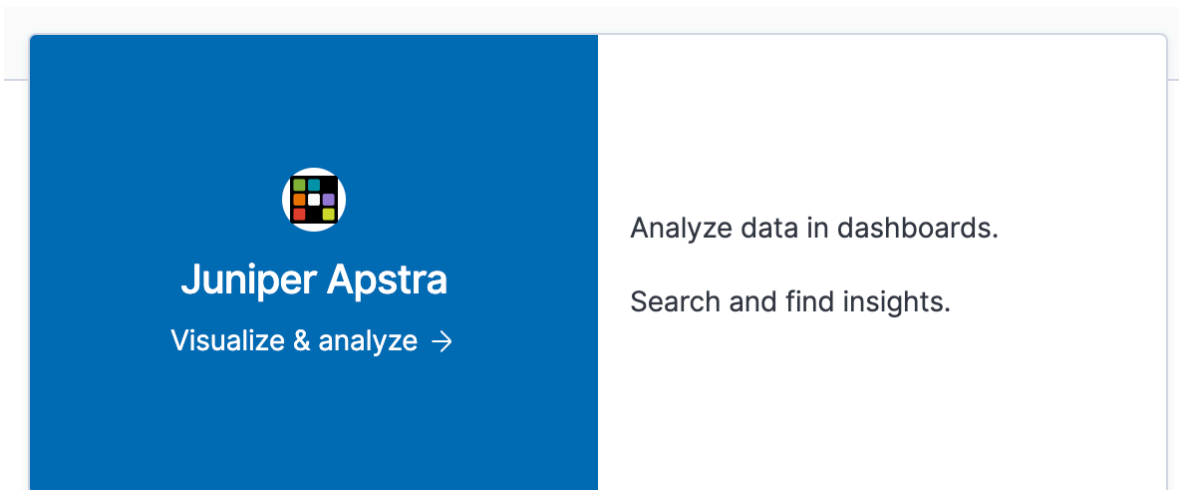
Global
The global tenant is shared between every OpenSearch Dashboards user.

Private
The private tenant is exclusive to each user and can't be shared. You might use the private tenant for exploratory work.

Choose from custom

[Cancel](#) [Confirm](#)

6. Click **Visualize and Analyze**.



7. Click **Dashboard** from the **OpenSearch Dashboards** page.

OpenSearch Dashboards [Add data](#)

Dashboard
Analyze data in dashboards.

Discover
Search and find insights.

Ingest your data

Add sample data
Get started with sample data, visualizations, and dashboards.

Manage your data

Interact with the OpenSearch API
Skip cURL and use a JSON interface to work with your data in Console.

[Make this my landing page](#) [View app directory](#)

8. Search for **Overview** in the Dashboards search bar.

Dashboards [Create Dashboard](#)

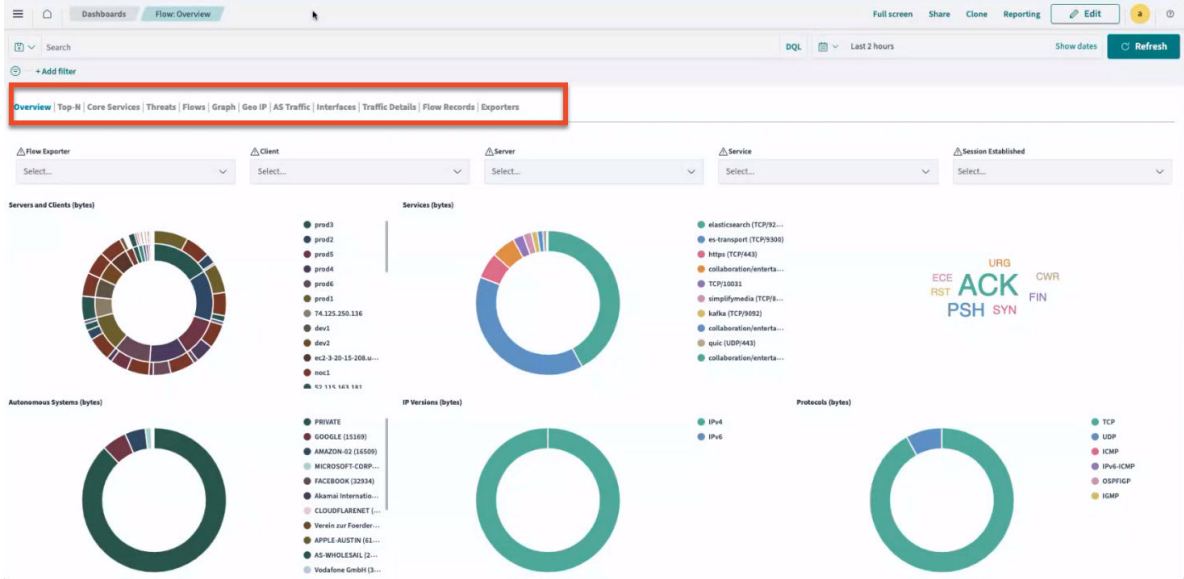
overf

<input type="checkbox"/>	Title	Type	Description	Last updated	Actions
<input type="checkbox"/>	Flow: Overview	Dashboard		Sep 27, 2023 @ 15:16:55.270	
<input type="checkbox"/>	Flow: Interfaces (overview)	Dashboard		Sep 27, 2023 @ 15:16:55.270	

Rows per page: 20

< 1 >

9. Select **Flow: Overview** to see the main dashboard. From here, you can view the various dashboards to analyze and visualize your data. Each dashboard shows a different perspective of your network traffic. For example:



Access the Apstra Flow Dashboard from the Apstra GUI

Although Apstra Flow has its own dashboard, you can also link directly to the dashboard from the Apstra GUI.

To access the Apstra Flow dashboard:

1. From the left navigation menu in the Apstra GUI, select **Analytics > Flow Data > Flow Collector**.
2. Enter the collector name, address, and username, then click **Create**.

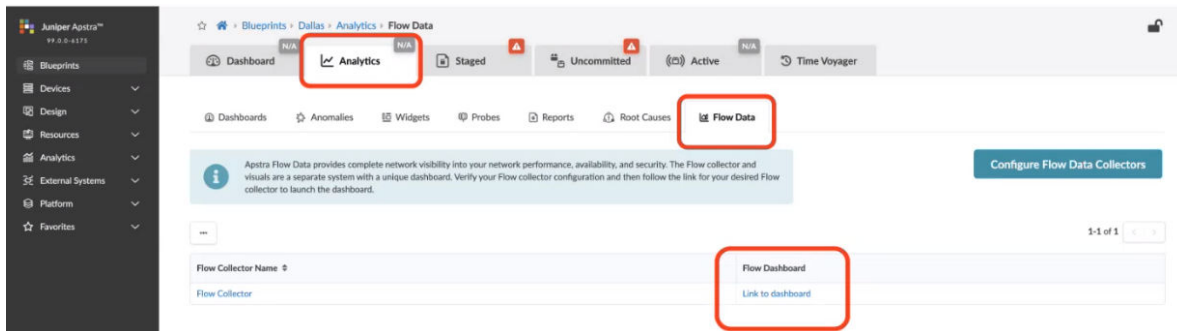
This information tells Apstra where your VM is located so you can access the flow collector and visualization dashboards.

The screenshot shows the 'Create Flow Collector' form. The fields are: Name (Flow), Address (10.28.36.6), Username (apstra), and Password (*****). There is a 'Create' button at the bottom right and a 'Create Another?' checkbox.

The screenshot shows the 'Flow Collectors' table with the following data:

Name	Address	Username	Actions
Flow	10.28.36.6	apstra	[Edit] [Delete]

- From your blueprint, navigate to **Analytics > Flow Data**. Click **Link to dashboard** to link directly to the Flow dashboard.



For more details about Apstra Flow, see the Analytics chapter in the [Juniper Apstra User Guide](#).

5

CHAPTER

Upgrade Apstra Flow

[Upgrade Apstra Flow on the Same VM \(In-Place\) | 28](#)

Upgrade Apstra Flow on the Same VM (In-Place)

SUMMARY

This topic provides instructions on how to upgrade your Apstra Flow configuration on the same VM. This process is referred to as an "in-place" upgrade.

IN THIS SECTION

- [Upgrade Apstra Flow | 28](#)

Upgrade Apstra Flow

You can run an automated upgrade script (recommended) to easily upgrade your Apstra Flow configuration with minimal downtime. Follow the steps in this procedure to upgrade Apstra Flow from version 4.2.1 to 5.0.0.

NOTE: Make sure that you upgrade to the latest Apstra version before running the upgrade file. You can find the current version by navigating to **Platform > About** in the Apstra GUI. The Apstra version is also shown in the left navigation menu of the GUI under the Juniper Apstra logo. For step-by-step upgrade instructions, see the [Juniper Installation and Upgrade Guide](#).

To automatically upgrade Apstra Flow:

1. Download the Apstra Flow upgrade file (apstra-flow-5.0.0-ubuntu-22.04-7.2.1-upgrade.tar.gz) from the [Juniper Support Downloads](#) page located under Applications.
2. Log in to the Apstra Flow console. The default credentials are: user=apstra and password=apstra.
3. Verify the flow collector version that is currently installed (in this case, 6.4.2).

```
apstra@apstra-flow:~$ /usr/share/juniper/bin/flowcoll -v
6.4.2.
```

4. Extract the upgrade (tar.gz) file you just downloaded in [Step 1](#). For example:


```
apstra@apstra-flow:~$ tar -xvzf apstra-flow-5.0.0-ubuntu-22.04-7.2.1-upgrade.tar.gz
```
5. When the files have finished downloading, cd to the /apstra-flow-5.0.0-ubuntu-22.04-7.2.1-upgrade directory.

```
apstra@apstra-flow:~$ cd apstra-flow-5.0.0-ubuntu-22.04-7.2.1-upgrade/
apstra@apstra-flow:~/apstra-flow-5.0.0-ubuntu-22.04-7.2.1-upgrade$ ls
advanced_settings.ndjson  flow-collector_7.2.1_linux_amd64.deb  opensearch-dashboards_2.16.0_amd64.deb  upgrade.sh
app.log                   opensearch_2.16.0_amd64.deb           README.md
dashboards.ndjson        opensearch-dashboards                 startup
apstra@apstra-flow:~/apstra-flow-5.0.0-ubuntu-22.04-7.2.1-upgrade$
```

NOTE: This folder contains a README.md file that includes steps on how to run a manual installation, if desired.

- Run the automated upgrade script. Accept the default settings when prompted during the installation. The upgrade will take around 3 to 5 minutes.

```
sudo ./upgrade.sh
```

- When the installation is complete, verify that the flow collector was upgraded to version 7.2.1.

```
apstra@apstra-flow:~/apstra-flow-5.0.0-ubuntu-22.04-7.2.1-upgrade$ /usr/share/juniper/bin/
flowcoll -v
7.2.1
```

- Check that your services are active and running.

Run the following commands to see the status of the flowcollector, opensearch, and opensearch-dashboard services:

```
sudo systemctl status flowcoll.service
```

```
apstra@apstra-flow:~/apstra-upgrade$ sudo systemctl status flowcoll.service
● flowcoll.service - Flow Collector
   Loaded: loaded (/etc/systemd/system/flowcoll.service; enabled; vendor preset: enabled)
   Drop-In: /etc/systemd/system/flowcoll.service.d
           └─flowcoll.conf
   Active: active (running) since Fri 2024-08-23 19:30:57 UTC; 1min 37s ago
   Main PID: 4715 (flowcoll)
     Tasks: 10 (limit: 24285)
    Memory: 350.6M
         CPU: 5.850s
   CGroup: /system.slice/flowcoll.service
           └─4715 /usr/share/juniper/bin/flowcoll --config /etc/juniper/flowcoll.yml

Aug 23 19:32:21 apstra-flow flowcoll[4715]: 2024-08-23T19:32:21.884Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:29 apstra-flow flowcoll[4715]: 2024-08-23T19:32:29.920Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:30 apstra-flow flowcoll[4715]: 2024-08-23T19:32:30.421Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:30 apstra-flow flowcoll[4715]: 2024-08-23T19:32:30.421Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:30 apstra-flow flowcoll[4715]: 2024-08-23T19:32:30.421Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:30 apstra-flow flowcoll[4715]: 2024-08-23T19:32:30.421Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:30 apstra-flow flowcoll[4715]: 2024-08-23T19:32:30.421Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:30 apstra-flow flowcoll[4715]: 2024-08-23T19:32:30.421Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:30 apstra-flow flowcoll[4715]: 2024-08-23T19:32:30.421Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
Aug 23 19:32:30 apstra-flow flowcoll[4715]: 2024-08-23T19:32:30.421Z    info    bootstrapper[opensearch].component_template_manager[opensearch]
```

```
sudo systemctl status opensearch
```

```
apstra@apstra-flow:~/apstra-upgrade$ sudo systemctl status opensearch
● opensearch.service - OpenSearch
   Loaded: loaded (/lib/systemd/system/opensearch.service; enabled; vendor preset: enabled)
   Active: active (running) since Fri 2024-08-23 19:30:45 UTC; 1min 57s ago
     Docs: https://opensearch.org/
   Main PID: 4319 (java)
     Tasks: 93 (limit: 24285)
    Memory: 10.6G
         CPU: 3min 10.880s
   CGroup: /system.slice/opensearch.service
           └─4319 /usr/share/opensearch/jdk/bin/java -Xshare:auto -Dopensearch.networkaddress.cache.ttl=60 -Dopensearch

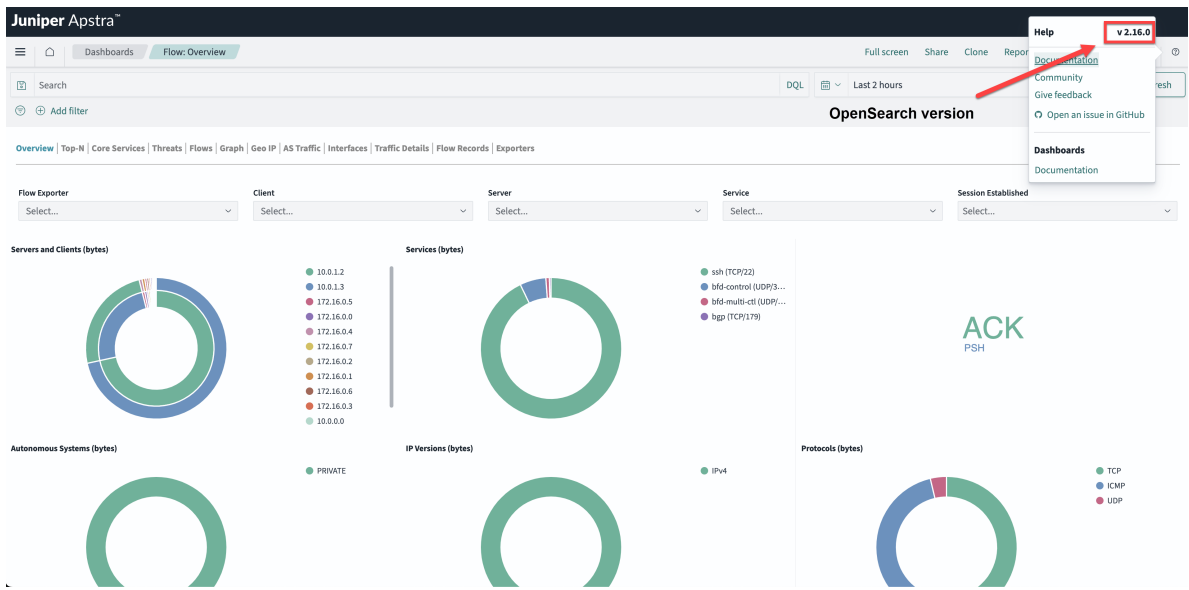
Aug 23 19:30:29 apstra-flow systemd-entrypoint[4319]: WARNING: System::setSecurityManager has been called by org.opensearch
Aug 23 19:30:29 apstra-flow systemd-entrypoint[4319]: WARNING: Please consider reporting this to the maintainers of org.op
Aug 23 19:30:29 apstra-flow systemd-entrypoint[4319]: WARNING: System::setSecurityManager will be removed in a future rel
Aug 23 19:30:30 apstra-flow systemd-entrypoint[4319]: Aug 23, 2024 7:30:30 PM sun.util.locale.provider.LocaleProviderAdapt
Aug 23 19:30:30 apstra-flow systemd-entrypoint[4319]: WARNING: COMPAT locale provider will be removed in a future release
Aug 23 19:30:30 apstra-flow systemd-entrypoint[4319]: WARNING: A terminally deprecated method in java.lang.System has been
Aug 23 19:30:30 apstra-flow systemd-entrypoint[4319]: WARNING: System::setSecurityManager has been called by org.opensearch
Aug 23 19:30:30 apstra-flow systemd-entrypoint[4319]: WARNING: Please consider reporting this to the maintainers of org.op
Aug 23 19:30:30 apstra-flow systemd-entrypoint[4319]: WARNING: System::setSecurityManager will be removed in a future rel
Aug 23 19:30:45 apstra-flow systemd[1]: Started OpenSearch.
```

sudo systemctl status opensearch-dashboards.service

```
apstra@apstra-flow:~/apstra-upgrade$ sudo systemctl status opensearch-dashboards.service
● opensearch-dashboards.service - "OpenSearch Dashboards"
   Loaded: loaded (/lib/systemd/system/opensearch-dashboards.service; enabled; vendor preset: enabled)
   Active: active (running) since Fri 2024-08-23 19:30:51 UTC; 2min 1s ago
     Main PID: 4605 (node)
        Tasks: 11 (limit: 24285)
       Memory: 175.2M
          CPU: 13.593s
   CGroup: /system.slice/opensearch-dashboards.service
           └─4605 /usr/share/opensearch-dashboards/node/bin/node /usr/share/opensearch-dashboards/src/cli/dist

Aug 23 19:30:58 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
Aug 23 19:30:58 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
Aug 23 19:30:58 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
Aug 23 19:30:58 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
Aug 23 19:30:58 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
Aug 23 19:30:58 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
Aug 23 19:30:58 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
Aug 23 19:30:58 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
Aug 23 19:31:22 apstra-flow opensearch-dashboards[4605]: {"type":"log","@timestamp":"2024-08-23T19:31:22Z","tags":["lis
Aug 23 19:31:57 apstra-flow opensearch-dashboards[4605]: [agentkeepalive:deprecated] options.freeSocketKeepAliveTimeout
```

9. Finally, reopen the Aspra Flow dashboards from the Apstra GUI and verify that OpenSearch was updated to v2.16.0.



Congratulations! You have successfully upgraded your Apstra Flow configuration. You can now start using the new features for Apstra 5.0.0.