Engineering
Simplicity

Juniper Mist Wired Assurance Configuration Guide

Published
2026-01-29

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

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 Mist Wired Assurance Configuration Guide

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

6 Troubleshooting

7 Appendix

About This Guide

Welcome to the Juniper Mist Wired Assurance Configuration Guide. This document is designed to assist network administrators and IT professionals in configuring Juniper switches using the Mist portal. The Mist portal provides a robust set of features for switch configuration and management, all delivered through the Juniper Mist Assurance cloud service.

To configure switches through the Mist portal, you must be assigned the Mist Super User role. This role grants the necessary permissions to perform switch configuration and management tasks within the portal. Ensure your account has the appropriate access before proceeding with any configuration steps.

1

CHAPTER

Get Started

IN THIS CHAPTER

Juniper Mist™ makes it easy to get started with your wired network. Use the information and links in this chapter to ease your way into the process of selecting hardware, configuring devices, tracking performance, and assessing user experiences.

Juniper Mist Wired Assurance Overview

SUMMARY

Watch these videos to start getting familiar with Juniper Mist™ Wired Assurance.

Juniper Mist™ Wired Assurance is an AI-driven cloud service that brings cloud management and Mist AI to enterprise campus switches. Wired Assurance simplifies all aspects of switch management that include device onboarding, configuration at scale, and monitoring and troubleshooting.

With Wired Assurance, you get real-time visibility into the health and performance of your wired network. You can see how your switches are doing, check out service level expectations (SLE) metrics, and even get insights into the end user experiences.

For a quick overview of Wired Assurance, watch the following video:

Video: Mist Wired Assurance Overview

When it comes to switch configuration, Wired Assurance lets you use configuration templates to easily apply consistent configurations across all your sites and devices, providing a streamlined switch management experience. Wired Assurance also has handy tools and features that help you troubleshoot network issues easily.

Wired Assurance is available as a subscription-based service right through the Juniper Mist portal.

Wired Assurance supports EX and QFX Series switches. We recommend using EX Series switches in places where you need interoperability with Juniper Mist Access Points (APs). To find out which switches are supported by Juniper Mist Wired Assurance, refer to Juniper Mist Supported Hardware.

Watch the following video to understand how Wired Assurance can automate and simplify device provisioning, deployment, and operation.

Video: Wired Assurance - Day 0, Day 1, and Day 2+

RELATED DOCUMENTATION

Overview of Template-Based Switch Configuration | 15

Hardware and Software for Your Wired Network

SUMMARY

Be aware of this hardware and software information as you get started installing, onboarding, and configuring your devices.

IN THIS SECTION

Hardware

Juniper provides a wide range of hardware to support your wired networking needs. Juniper Mist Wired Assurance supports onboarding, configuration, and management of Juniper EX Series and QFX Series switches. Click the links below to access datasheets, quick start guides, and hardware guides for these switches on the Juniper Mist Supported Hardware page:

- EX Series Switches
- QFX Series Switches

NOTE: The information provided on the Juniper Mist Supported Hardware page is not specific to Wired Assurance. This page lists all the devices that can be managed through the Mist portal.

Firmware and Management Options

- Minimum OS Release—The minimum required Junos OS release (firmware image) for Juniper switches across platforms is 18.2R3. Be aware that 18.2R3 has reached end of support. We recommend that you upgrade to a JTAC-suggested Junos release. For the suggested releases, refer to Junos Software Versions – Suggested Releases to Consider and Evaluate. If you have any questions, write to support@mist.com.

- FIP Mode Not Supported—Wired Assurance does not support the use of the FIPS mode command or management of devices that are already in FIPS mode.

Switch Administrator Role Requirements

SUMMARY

Compare job tasks to admin roles to ensure that your admin users have the appropriate privileges.

Before you onboard and configure your switches, ensure that you have the required switch administrator role.

The following table lists the available privileges for each switch administrator role (Super User, Network Admin, Help Desk, and Observer). A check mark next to a privilege means that the user role enjoys that privilege. An x means that the user role does not enjoy that privilege.

Privileges	Super User	Network Admin (All Sites Access)	Network Admin (Site Group or Specific Sites Access)	Helpdesk	Observer
Claim switches	✓	×	×	×	×
Adopt switches	✓	✓	✓	✓	✓
Release switches	✓	×	×	×	×
View switch details	✓	✓	✓	✓	✓
Access utility tools (ping, traceroute, cable test, bounce port)	✓	✓	✓	×	×
Access switch shell	✓	✓	✓	×	×

(Continued)

Privileges	Super User	Network Admin (All Sites Access)	Network Admin (Site Group or Specific Sites Access)	Helpdesk	Observer
Reboot the switch	✓	✓	✓	×	×
Edit, save, and apply switch configuration from the **Switches** page or the **Site** > **Switch Configuration** page.	✓	✓	✓	×	×
Access switch template	✓	×	×	×	×
Assign switch template to sites	✓	✓	×	✓ (Applicable only to roles with access to all sites. Not available to roles with access to all site groups or specific sites.)	✓ (Applicable only to roles with access to all sites. Not available to roles with access to all site groups or specific sites.)
Enable/disable switch configuration management	✓	✓	✓	×	×
Send the switch logs to the Mist cloud	✓	✓	✓	×	×
View the Inventory page	✓	✓	×	×	✓

(Continued)

Privileges	Super User	Network Admin (All Sites Access)	Network Admin (Site Group or Specific Sites Access)	Helpdesk	Observer
Assign to a site	✓	✓ (Applicable only to the site assignment option on the switch details page)	×	×	×
Rename the device	✓	✓ (Applicable only to the rename option on the switch details page)	✓	×	×
Access the switch management root password	✓	✓	×	×	×
Access to wired SLE, wired clients, the wired insights switch, or wired insights clients	✓	✓	✓	✓	✓

Wired Network Deployment Workflow

SUMMARY

Complete these essential tasks to set up your organization and sites, ensure security, install your devices, and start configuring your network.

Table 1: Deployment Tasks and Links

Category	Task	More Information
Prerequisites	Before you can configure your wired network or onboard your devices, you need to complete these tasks in the Juniper Mist™ portal: • Create your organization, set up at least one site, and activate your subscriptions. • Add user accounts for other personnel who are working with you to deploy Juniper Mist. You can even enable limited access for the personnel who are installing devices. • Configure your firewall to allow Juniper Mist traffic. • Set up other security options as needed. For example, manage certificates, disable Juniper Mist support access, or enable Single Sign-On.	• Juniper Mist Quick Start • Firewall Configuration: Juniper Mist IP Addresses and Ports • Security Options

Table 1: Deployment Tasks and Links ***(Continued)***

Category	Task	More Information
Understand Admin Permissions	Make sure that your admin account gives you the permissions that you need for your configuration tasks.	"Switch Administrator Role Requirements" on page 4
Onboard Switches	Add switches to your Juniper Mist organization, either in a greenfield (new cloud-ready switches) or a brownfield (previously deployed) approach.	"Onboard Switches to Mist Cloud" on page 16
Configure Switch Templates	For large-scale deployments, instead of configuring each switch individually, we recommend using switch configuration templates to set up and streamline configurations across multiple sites. For smaller deployments or special use cases, you can also configure at the site-level and device-level.	"Create a Switch Configuration Template" on page 22
Assign the Templates to Site	After creating a switch configuration template, assign it to the relevant sites.	"Assign a Template to Sites" on page 22
Configure Site-specific Settings	This step is applicable if you want to make site-level adjustments to the switch configuration.	"Configure Site-Specific Settings" on page 23
Configure Switch-specific Settings	In an organization level deployment, you might also need to configure parameters that are specific to each switch and cannot be configured via the template. Also, for specific use cases, you can configure an entire switch individually.	"Configure Switch-Specific Settings" on page 24

Table 1: Deployment Tasks and Links ***(Continued)***

Category	Task	More Information
Verify Switch Configuration	Review the configuration applied to your switches and make any updates through the switch details page.	"Verify Switch Configuration" on page 27
Build Campus Fabric	After switches are onboarded and configured, you can include them in your campus fabrics topologies.	• "Configure Campus Fabric EVPN Multihoming" on page 257 • "Configure Campus Fabric Core-Distribution" on page 265 • "Configure Campus Fabric IP Clos" on page 275

Request Help with a New Deployment

SUMMARY

If you need help with a new deployment of Juniper Mist™ devices, follow these steps to submit your request.

The Juniper Mist Support team provides onboarding assistance to help customers with new deployments. You'll get assistance with initial setup, configuration, and basic troubleshooting.

NOTE:

- These services do not include network design.
- Submit your request at least 48 hours in advance of your preferred appointment time.

- Available only for wireless, switching, and SD-WAN deployments.

To request help with a new deployment:

1. Click the question icon near the top-right corner of the Juniper Mist portal, and then click **Support Tickets**.

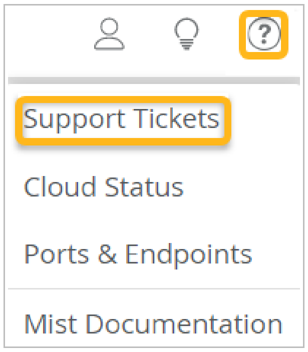

2. Click **Create a Ticket**.
3. For **Technology**, select the technologies that are applicable to your network and are included in the onboarding support program.

> **NOTE**: Onboarding help is available only for wireless, switching, and SD-WAN deployments.

4. For **Ticket Type**, select **Onboarding Help for New Deployment**.

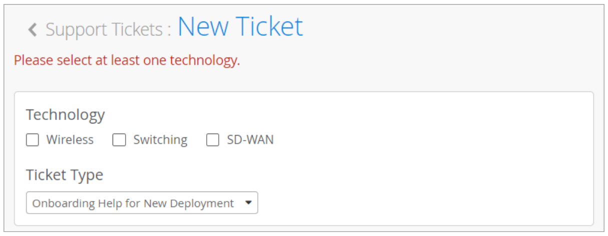

< Support Tickets : New Ticket

Please select at least one technology.

Technology

□ Wireless □ Switching □ SD-WAN

Ticket Type

Onboarding Help for New Deployment ▾

5. Enter a short **Ticket Summary** and a detailed **Description**.
6. Under **Checklist**, answer all the questions.

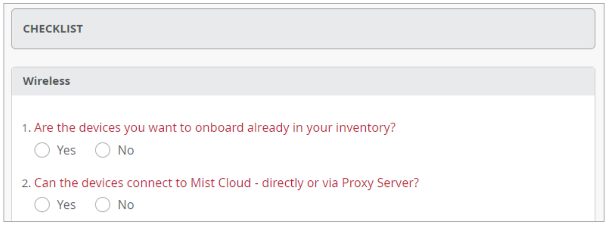

CHECKLIST

Wireless

1. Are the devices you want to onboard already in your inventory?
○ Yes ○ No

2. Can the devices connect to Mist Cloud - directly or via Proxy Server?
○ Yes ○ No

> **NOTE**: The checklist includes each technology that you selected at the top. Complete all questions in all sections that appear.
>
> As you complete the checklist, you can use suggested hyperlinks for self-help. If you find your answers in these links, you can cancel submitting this ticket.

7. Under **Schedule**, select the date, time, and time zone for your onboarding help session.

> **NOTE**: If you need to change your appointment after you submit your ticket, you can go to your list of open tickets, select this ticket, and reschedule.

8. Click **Submit** at the top right corner of the page.

   If this button is unavailable, ensure that you've entered all required information. Required fields have red labels.

The support team will contact you to conduct the help session.

Explore Juniper Mist Features

SUMMARY

Get familiar with a few features that you might find especially useful.

Now that your wired network is up and running, explore other Juniper Mist™ features to meet your business needs.

Here are some features we think you'll find especially helpful.

- Switch Dashboard—Track the switch performance against compliance parameters. See:
  - "Switch Metrics" on page 293
  - "Switch Details" on page 294
  - "Switch Utilities" on page 307
- Wired Service Level Expectations (SLEs)—Use the SLE dashboards to assess the network's user experience and resolve any issues proactively. See "Wired SLEs" on page 313.

- Port Profiles—Port profiles provide a convenient way to manually or automatically provision switch interfaces. See "Port Profiles" on page 71.
- Campus Fabric—Juniper Networks campus fabrics provide a single, standards-based Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) solution that you can deploy on any campus. You can deploy campus fabrics on a two-tier network with a collapsed core or a campus-wide system that involves multiple buildings with separate distribution and core layers. To get started, see "Determine a Campus Fabric Topology" on page 245.
- Group Based Policy—A group-based policy (GBP) helps you achieve microsegmentation and macrosegmentation, for example to secure data and assets, in Virtual extensible Local Area Network (VXLAN) architecture. See "Configure BGP on Switches via Mist" on page 89.
- Virtual Chassis—The Virtual Chassis technology enables you to connect multiple individual switches together to form one logical unit and to manage the individual switches as a single unit. See "Virtual Chassis Overview (Juniper Mist)" on page 208.

2

CHAPTER

Switch Configuration

IN THIS CHAPTER

Overview of Template-Based Switch Configuration

SUMMARY

Before you configure your switches, get familiar with configuration templates and the hierarchical configuration model in Juniper Mist™.

The Mist portal is a handy tool that simplifies the whole switch configuration process. One of its cool features is the template-based, hierarchical configuration model. Instead of configuring each switch individually, you can use a configuration template to set up and streamline configurations across multiple sites. See "Create a Switch Configuration Template" on page 22 for more information on how to configure switches.

Any device connected to a particular site will inherit the template settings applied to that site. The configuration inheritance model follows this hierarchy: organization-level template > site-level configuration > device-level configuration. In this hierarchy, the organization-level template provides the global settings that are applied to all the switches managed by it. Any site-specific updates will apply to all the devices in a site. You can configure any device-specific configuration updates (such as, adding hostname, switch role, and IRB interfaces) at the individual switch-level.

When a conflict between the organization-level template settings and site-level configuration settings occurs, the narrower settings override the broader settings. For example, when settings at both the template and site levels apply to the same device, the narrower settings (in this case, site settings) override the broader settings defined at the template or organization level.

The configuration template also has options to include additional CLI commands in the `set` format to configure settings, for which the template doesn't provide GUI options. When additional CLI commands are included in site-level configuration, they are added using an OR logic. If you include additional CLI commands in the device-specific configuration, they are included using an AND operation, where all the CLI commands are applied together.

Also, you can use the port configuration feature in the organization template to create different port configuration rules for each of the switch models found in the organization. For more information, see "Port Profiles" on page 71.

To further simplify your configuration tasks, Mist also provides an option to use site variables to streamline the switch configuration. Site variables, configured at **Organization** > **Site Configuration** > **Site Variables**, provide a way to use tags to represent real values so that the value can vary according to the context where you use the variable. This means the same variable can configure different values in different sites. The fields that support configuration through site variable have a help text showing the site variable configuration format underneath them.

To configure site variables, follow the steps provided in Configure Site Variables.

Onboard Switches to Mist Cloud

SUMMARY

Follow these steps to claim new switches or adopt previously deployed switches into your organization.

IN THIS SECTION

To configure and manage a switch through Juniper Mist cloud, you must onboard the switch into your organization.

NOTE:

- Wired Assurance does not support Junos Flex images. To ensure compatibility, please verify that your switch is running a standard (non-Flex) Junos image. When upgrading a switch, we recommend doing so via the Mist cloud, which ensures that only a standard Junos image is deployed.
- Ignore the steps in this topic if your switches are already onboarded to the Mist cloud.

You can onboard greenfield or brownfield switches to Mist:

- Greenfield—New cloud-ready switches. Cloud-ready devices have a QR code or claim code on the chassis that you can scan to onboard the device quickly in the Juniper Mist portal.
- Brownfield—Existing, in-service (brownfield) switches which are not cloud-ready but are still supported by Mist. These switches do not have a QR code attached. For devices that are not cloud-ready but still supported in Mist, you must enter some Junos CLI commands locally in the switch to onboard (adopt) the device.

Juniper switches use an SSH connection (**TCP port 2200)** to send telemetry updates to the Juniper Mist cloud. From the Juniper Mist portal, SSH is also used when adopting a brownfield switch, and when opening a remote shell to the switch from the Utilities menu. In addition, if you include any custom CLI configurations in the Additional CLI commands section of the Switch configuration page, these commands are pushed to the switch via SSH.

To forestall the chance of denial-of-service attacks or other possible SSH exploits, you can configure an upper limit of connections and/or sessions per connection on the switch. For Juniper EX and QFX series switches, the default value for *ssh max-sessions-per-connection* is 10, You can change that by issuing an Junos command such as the following:

```
[edit system services ssh] set connection-limit <number>
```

Switch Onboarding Prerequisites

Before you onboard a switch:

- Ensure that you have a Wired Assurance Subscription, and login credentials for the Juniper Mist portal. To get started with Mist, follow the instructions in Quick Start: Mist.
- Ensure that the switch is supported by Mist. To see the switch models supported by Mist, visit Juniper Mist Supported Hardware.
- Ensure that the switch is connected to a DNS server (an NTP server is also recommended), and is able to connect to the Juniper Mist cloud architecture over the Internet.
- If there is a firewall between the cloud and the switch, allow outbound access on TCP port 2200 to the management port of the switch.

Onboard a Greenfield Switch

You can onboard a single greenfield, cloud-ready switch to the Mist cloud via the Mist AI Mobile App. However, if you want to onboard multiple cloud-ready switches together, you can do that via the Juniper Mist portal, by using the activation code associated with the purchase order.

To onboard a greenfield switch, follow the instructions in Quick Start: Cloud-Ready EX and QFX Switches with Mist.

For a quick demo, watch the following video:

Video: Onboard One or More Switches Using a Web Browser

Onboard a Brownfield Switch

Use the **Adopt Switch** option to onboard a brownfield switch that is not cloud-ready or does not have a QR code. When you adopt a brownfield switch, you can have Mist manage it (recommended), which means any existing configuration will be replaced with settings made in the Mist console. Or, you can choose to *not* have Mist manage the switch, in which case the existing configuration will remain as is, plus some new settings for connecting to the cloud and telemetry. An unmanaged switch will be unique – configurations made in the Mist console will not be applied, nor will it benefit from the use of templates or site variables or any of the other conveniences available to managed devices. In addition, subsequent configurations made on the switch will not be "known" to Mist, so you may want to set up a warning message in the CLI to indicate it is part of the Mist environment, or limit who can make configuration changes so they don't conflict.

If Mist will manage the switch, that is, you select the **Manage configuration with Mist** option when onboarding it, be sure to back up the existing Junos OS configuration before adopting the switch. Do this by connecting to the switch, logging on to the CLI, and in Junos, running the `request system configuration rescue save` command to save the currently active configuration and any installation-specific parameters.

For virtual devices such as a vJunos-switch or legacy devices that predate the use of Mist claim codes on the hardware, you need to adopt, rather than claim them. Note that if the VM was previously claimed in one environment, such as Global02, it may not be available from the inventory or installed base of another environment or organization (this is because the MAC address may still be attached to the original environment.) You need to release the device from the original environment, or recreate the virtual device, which will generate a new virtual MAC for it.

In the procedure below, you will make configuration changes to the Juniper Mist portal, and also to the switch using the Junos OS CLI. Be sure you can log in to both environments.

NOTE: An unmanaged switch still receives configurations from Mist to maintain connectivity with the Mist cloud. These include system scripts and extensions for efficient stats collection, system syslog settings for efficient logging on the device, and a user account named 'mist' for cloud communication.

To onboard a brownfield switch to the Mist cloud:

1. Log in to your organization on the Juniper Mist cloud and then click **Organization** > **Inventory** in the menu.
2. Select **Switches** at the top of the page that appears, and then click the **Adopt Switch** button in the upper-right corner to generate the Junos OS CLI commands needed for the interoperability. The commands create a Juniper Mist user account, and an SSH connection to the Juniper Mist cloud over TCP port 2200 (the switch connection is initiated from a management interface and is used for setting up configuration and sending telemetry data).

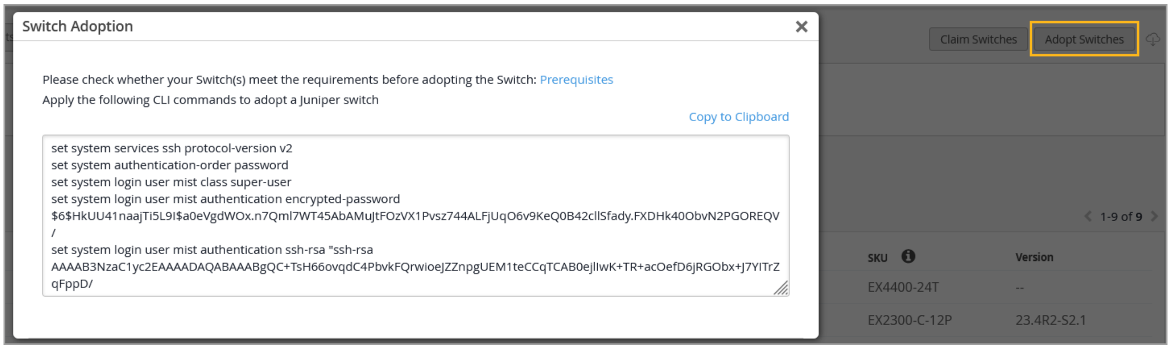

Please check whether your Switch(s) meet the requirements before adopting the Switch: Prerequisites

Apply the following CLI commands to adopt a Juniper switch

```
set system services ssh protocol-version v2
set system authentication-order password
set system login user mist class super-user
set system login user mist authentication encrypted-password
$6$HkUU41naajTi5L9I$a0eVgdWOx.n7Qml7WT45AbAMuJtFOzVX1Pvsz744ALFjUqO6v9KeQ0B42cllSfady.FXDHk40ObvN2PGOREQV
/
set system login user mist authentication ssh-rsa "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABgQC+TsH66ovqdC4PbvkFQrwioeJZZnpgUEM1teCCqTCAB0ejllwK+TR+acOefD6jRGObx+J7YITrZ
qFppD/
```

3. In the page that appears, click **Copy to Clipboard** to get the commands from the Juniper Mist cloud.
4. Log in to the switch via Junos OS CLI.
5. In the CLI, type `edit` to start configuration mode, and then paste the commands you just copied (type `top` if you are not already at the base level of the hierarchy).
6. (Optional) If you want to add a system message, use the following command:

```
user@switch# set system login message message text here
```

7. You can confirm your updates on the switch by running `show` commands at the `[system services]` level of the hierarchy, and again at the `[system login user mist]` level of the hierarchy.

```
user@switch# show system services
ssh {
    protocol-version v2;
}
netconf {
    ssh;
}
outbound-ssh {
    client juniper-mist {
        device-id 550604ec-12df-446c-b9b0-eada61808414;
        secret "trimmed"; ## SECRET-DATA
        keep-alive {
            retry 3;
            timeout 5;
        }
        services netconf;
        oc-term.mistsys.net {
            port 2200;
            retry 1000;
            timeout 60;
        }
    }
}
dhcp-local-server {
    group guest {
        interface irb.188;
    }
    group employee {
        interface irb.189;
    }
    group management {
        interface irb.180;
    }
}
```

```
user@switch-1#  show system login user mist
class super-user;
authentication {
    encrypted-password "$trimmed ## SECRET-DATA
}
```

8. Run the `commit` command to save the configuration.
9. On the Juniper Mist portal, click **Organization** > **Inventory** > **Switches** and select the switch you just added.
10. Click the **More** drop-down list at the top of the page, and then click the **Assign to Site** button.
11. In the page that appears, choose which site you want to assign the switch to.
12. Select the **Manage configuration with Mist** check box.
13. Click **Click Assign to Site**.

    The onboarded switch, when assigned to a site, will appear on the "Switches (List View)" on page 291 page with the status Connected.

    You can also run the `show system connections` command locally on the switch to check whether the switch is connected to Mist. This command shows an **ESTABLISHED** TCP session to Mist.

For a quick demo, watch the following video:

Video: Onboard a Brownfield Switch

Configure Switches Using Templates

SUMMARY

Follow these steps to use templates to configure your switches with the exact features and settings that you need for the users at your sites.

IN THIS SECTION

We recommend that all switches in an organization be managed exclusively through the Juniper Mist cloud, and not from the device's CLI.

The recommended approach to configuring multiple switches with Juniper Mist™ Wired Assurance involves two main steps: creating a switch configuration template and applying it to one or multiple sites. The configuration settings linked to a particular site will be applied to the switches (both newly added and existing switches) within that site. This allows you to manage and apply consistent and standardized configurations across your network infrastructure, making the configuration process more efficient and streamlined.

You can also configure switches individually (without using the template) from the "Switch Details" on page 294 page.

For a quick overview of the switch templates, watch the following video:

Video: Configuration Models (Global Templates)

To configure a switch, you need to have a Super User role assigned to you. This role grants you the necessary permissions to make changes and customize the switch settings.

To find out which switches are supported by Juniper Mist Wired Assurance, refer to Juniper Mist Supported Hardware.

Create a Switch Configuration Template

Switch configuration templates make it easy to apply the same settings to switches across your sites. Whether it's one site or multiple sites, you can use the template to quickly configure new switches. When you assign a switch to a site, it automatically adopts the configuration from the associated template.

> **NOTE**: Configuration done on the switch through the Mist dashboard overrides any configuration done through the device CLI. The switch details page doesn’t display any configuration changes you make directly on the switch through the switch CLI.

To create a switch configuration template:

1. Open the Juniper Mist™ portal and click **Organization** > **Switch Templates**.
2. Click **Create Template**, enter a name for the template in the **Template Name** field, and then click **Create**.

   The Switch Templates page appears. The selected template is identified at the top of the page.

   > **NOTE**: If you prefer, you can import the template settings in a JSON file instead of manually entering the information. To import the settings, click **Import Template**. To get a JSON file with the configuration settings that can be customized and imported, open an existing configuration template of your choice and click **Export**. For more information, refer to "Manage Templates Settings" on page 150.

3. Enter your settings.

   For more information, see "Switch Configuration Options" on page 29.
4. Click **Save** to save the switch template.

   The **Confirm changes** window appears.
5. Click **Save** on the **Confirm changes** window.

   The template is saved. To view the new template, go to **Organization** > **Switch Templates**.

Assign a Template to Sites

After creating a switch configuration template, you need to assign it to the relevant sites. This ensures that the configuration settings are applied to the devices within those sites. You have the flexibility to apply the template to a single site or multiple sites, depending on your specific requirements.

To assign a template to one or multiple sites:

1. Click **Organization** > **Switch Templates**.

   The **Switch Templates** page appears.
2. Click the template that you want to assign to sites.

   The **Switch Templates:** ***Template-Name*** page appears.
3. Click **Assign to Sites**.

   The **Assign Template to Sites** window appears.
4. Select the sites to which you want to apply the template and then click **Apply**.

Alternatively, you can apply a template to a site from the **Site Configuration** page, using the following steps:

1. Click **Site** > **Switch Configuration**.

2. Click a site from the list to open it.

3. Select a template from the **Configuration Template** field, and then click **Save**.

Configure Site-Specific Settings

Switches connected to a particular site inherit the organization-level template settings applied to that site. If required, you can customize the settings at the site level. These site-specific changes will not affect the organization-level template and will apply to all the switches in the selected site.

NOTE: Unless required, do not override the settings at the site level. Use site variables, instead.

To update switch settings at the site level:

1. Click **Site** > **Switch Configuration**.
2. Click a switch configuration line item to open it.

   The switch configuration will be populated with settings from the configuration template assigned to the selected site. If you want to assign a different template to the site, select it from the Configuration Template field.
3. To override the template settings applied to the site, follow these steps:

   a. Select the **Override Configuration Template** check box in relevant tiles.

   b. Edit the settings.

      For more information, see "Switch Configuration Options" on page 29.
4. After making the changes, click **Save**.

The **Confirm changes** window appears.

5. Click **Save** to apply the updated settings.

   The changes are applied at the site level.

Configure Switch-Specific Settings

IN THIS SECTION

You need to configure certain parameters on individual switches. This can be specific to each switch and cannot be configured via the template. The switch-specific settings could include a switch name, role, management interface (out of band), and an IRB interface. You can either configure the settings manually on the individual switches, or import the settings.

Configure Switch-Specific Settings Manually

To configure additional switch-level configuration settings manually:

1. Click **Switches** on the left menu.
2. From the switch list, click the switch you want to edit.
3. Configure the switch-specific settings that include the following:
   - INFO— Configure the details on the INFO tab. The details include a hostname for the switch and the role of the switch in the network (example: Access).

   - IP Configuration (Out of Band)—The management interface settings. If you want to override the template settings, select the **Override Site/Template Settings** check box, and then make the changes. You can specify if the IP address is static or DHCP-based. You can also enable or disable **Dedicated Management VRF** (out of band). For all standalone devices or Virtual Chassis running Junos version 21.4 or later, this feature confines the management interface to non-default virtual routing and forwarding (VRF) instances. Management traffic no longer has to share a routing table with other control traffic or protocol traffic.

   - IP Configuration—The IRB interface settings for inter-VLAN routing. If you want to override thes template settings, select the **Override Site/Template Settings** check box, and then make the changes. You can also add new IP configuration by clicking **Add IP Configuration**. You can specify if the IP address is static or DHCP-based. You can specify multiple IP addresses in the Additional IP Configuration section. Mist supports IPv4 and IPv6 addresses. The IPv6 support is available for IP address and subnet mask.

> **NOTE**: If the IP address specified in the Additional IP Configuration section under IP Configuration does not fall within the scope of the subnet configured in the associated network (VLAN), the IP Configuration window displays a warning message to indicate the mismatch.

- Port Configuration—Configure switch ports and apply port profiles to them. Layer 3 (L3) interface in Port Configuration support IPv4 and IPv6 addresses. The IPv6 support is available for IP address and subnet mask. You can add new port configuration by clicking **Add Port Configuration**. Other port configurations include the following:

  - Q-in-Q tunneling—For more information, refer to "Configure Q-in-Q Tunneling on a Switch Port" on page 148.

  - Port channelization and speed configuration—For more information, refer to "Configure Port Speed" on page 84.

- Bridge Priority—See "Configure Bridge Priority on Switches via Mist" on page 136.

- Additional CLI Commands—If you include any additional CLI commands in the device-specific configuration, they are included using an AND operation, where all the CLI commands are applied together.

- OSPF—You can enable OSPF configuration from the OSPF tile. When enabling OSPF, you can also configure a routing policy (import or export policy) to be applied to OSPF routes. For more information, refer to "OSPF Configuration for Switches" on page 137.

4. If you want to override the template settings applied to the switch, follow these steps:

   a. Select the **Override Site/Template Settings** check box in relevant tiles.

   b. Edit the settings and then click **Save**. The changes are immediately applied to the switch.

Configure Switch-Specific Settings Using the Bulk Upload Option

If you don't want to manually configure the switch-specific settings on each switch, you can configure the settings by uploading them through a CSV file. You can upload the settings for one or more switches at once. You can upload the following settings: MAC address, serial number, switch name, switch role, router-ID, IP configuration (OOB), Primary IP (In-Band), and Default Gateway (In-Band).

To upload the switch level settings:

1. Click **Switches** on the left menu.

2. From the switch list, select the switches you want to configure.

   You can select one or more switches. These switches can be in connected or disconnected state. You can select switches regardless of whether they have configuration management enabled in Mist or not. However, we recommend that you disable configuration management on the devices before you perform this configuration update, and enable it back after the switch configuration update is completed. This approach prevents any unwanted configuration overrides.

3. Click **Bulk Upload Configuration**. The Bulk Upload Configurations window appears.

   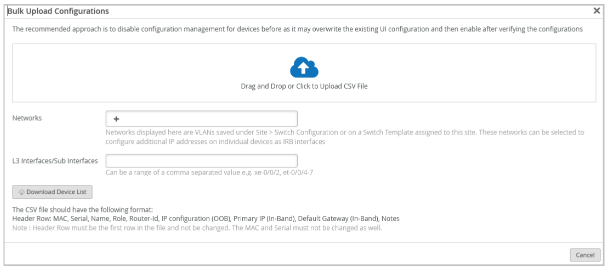

   Bulk Upload Configurations

   The recommended approach is to disable configuration management for devices before as it may overwrite the existing UI configuration and then enable after verifying the configurations

   Networks: +

   Networks displayed here are VLANs saved under Site > Switch Configuration or on a Switch Template assigned to this site. These networks can be selected to configure additional IP addresses on individual devices as IRB interfaces

   L3 Interfaces/Sub Interfaces:

   Can be a range of a comma separated value e.g. xe-0/0/2, et-0/0/4-7

   The CSV file should have the following format:
   Header Row: MAC, Serial, Name, Role, Router-Id, IP configuration (OOB), Primary IP (In-Band), Default Gateway (In-Band), Notes
   Note : Header Row must be the first row in the file and not be changed. The MAC and Serial must not be changed as well.

4. Download a sample CSV file from the Bulk Upload Configurations window by clicking **Download Device List**.

   NOTE:

   - If you don't need a sample file, you can use your own custom configuration file directly.

   - If you want any networks or L3 interfaces/sub Interfaces configuration to be present in the sample CSV downloaded, specify those on the Bulk Upload Configurations window before downloading the file. The downloaded sample file includes fields to configure settings for the specified networks and interfaces. The network selection allows you to configure additional IP addresses on individual devices as IRB interfaces.

   - You can add only one VLAN to an L3 sub interface. Only the networks created in the switch configuration or switch template can be added to the L3 sub interface configuration.

5. Update it with the required information in accordance with the guidelines provided in the sample sheet.

   NOTE:

- All fields except Name, IP Configuration (OOB), and Primary IP (In-Band) are optional. The header row must be the first row in the CSV file. Don't modify the MAC addresses and the serial numbers in the CSV file.

- If any field in the CSV file is left empty, the corresponding field on the switch configuration will be updated with a null value. This means any existing value for that field will be removed from the switch configuration.

6. After you update the configuration file, use the **Drag and Drop or Click to Upload CSV File** option to upload it.

   You can use the guidelines on the Bulk Upload Configurations window to perform the upload.

7. When you open the file to be uploaded from file upload window, the UI page loads the configuration in an editable format as shown below,

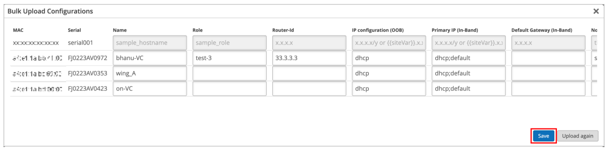

Bulk Upload Configurations

MAC	Serial	Name	Role	Router-Id	IP configuration (OOB)	Primary IP (In-Band)	Default Gateway (In-Band)	No
xx:xx:xx:xx:xx:xx	serial001	sample_hostname	sample_role	x.x.x.x	x.x.x.x/y or {{siteVar}}.x.:	x.x.x.x/y or {{siteVar}}.x.:	x.x.x.x	t
[redacted]	FJ0223AV0972	bhanu-VC	test-3	33.3.3.3	dhcp	dhcp;default		s
[redacted]	FJ0223AV0353	wing_A			dhcp	dhcp;default		
[redacted]	FJ0223AV0423	on-VC			dhcp	dhcp;default		

NOTE:

- If the CSV file does not contain information for some of the switches you selected, the configuration will not be pushed to those switches (the ones that are missing in the file).

- If the CSV file contains information for switches you haven't selected, the configuration will not be pushed to those switches either.

8. After making any further changes (if required), click **Save**.

   A confirmation message, indicating the number of devices updated, is displayed.

Verify Switch Configuration

You can easily review the configuration applied to your switches and make any updates through the "switch details" on page 294 page on the Mist portal.

To access the switch details page:

1. On the Mist portal, click the **Switches** tab on the left menu to open the Switches page.

2. On the **List** tab, click a switch to open the switch details page.

When the switch details page opens, you'll find yourself on the Front Panel tab. This tab gives you a comprehensive overview of the switch's port panel.

To check the configuration and status of a specific port, hover over that port in the front panel illustration. For instance, if you hover over port ge-0/0/45 in the following example, you'll see information indicating that a Mist AP is connected to that port. The displayed information also includes details about speed, power, the IP address, and more.

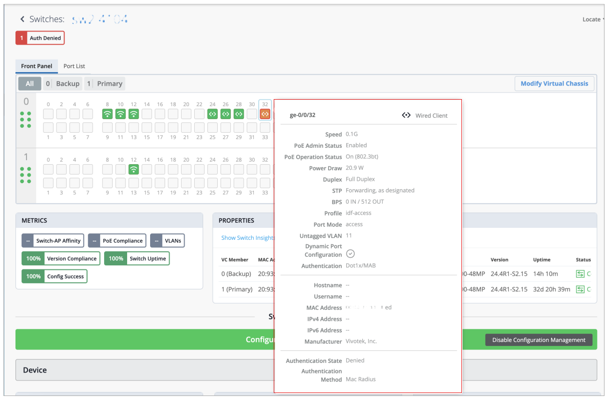

ge-0/0/32 Wired Client

Speed 0.1G
PoE Admin Status Enabled
PoE Operation Status On (802.3bt)
Power Draw 20.9 W
Duplex Full Duplex
STP Forwarding, as designated
BPS 0 IN / 512 OUT
Profile idf-access
Port Mode access
Untagged VLAN 11
Dynamic Port Configuration ✓
Authentication Dot1x/MAB

Hostname --
Username --
MAC Address [redacted]
IPv4 Address --
IPv6 Address --
Manufacturer Vivotek, Inc.

Authentication State Denied
Authentication Method Mac Radius

VC Member	MAC Ad		Version	Uptime	Status
0 (Backup)	20:93:	0-48MP	24.4R1-S2.15	14h 10m	C
1 (Primary)	20:93:	0-48MP	24.4R1-S2.15	32d 20h 39m	C

Click the port on the front panel illustration to see a more detailed view. From this view, you can perform tasks such as accessing the connected devices (for example, APs), viewing switch insights and editing the port configuration.

On the switch details page, you can also find information about switch events such as configuration changes in the "Switch Insights" on page 301 section.

If you want to download the configuration in a text file, select the **Download Junos Config** option on the **Utilities** drop-down list on the switch details page.

To see the complete configuration applied to the switch, simply scroll down to the **Switch Configuration** section. From there, you can view and, if needed, edit the configuration elements.

If required, you can update the settings at the switch level, site level, or template level. You can also use CLI commands to configure features that the predefined drop-down lists and text fields on the Mist portal do not support. For more information on how to update the settings, refer to "Manage or Update Configuration Settings" on page 149.

SEE ALSO

Manage or Update Configuration Settings | **149**

Switch Details | **294**

No Link Title

Switch Configuration Options

SUMMARY

Use this information to configure your switches.

IN THIS SECTION

Overview

You can enter switch settings at the organization level or the site level.

- To configure organization-wide settings, select **Organization** > **Switch Templates** from the left menu of the Juniper Mist portal. Then create your template and apply it to one or more sites or site groups.
- To configure switch settings at the site level, select **Site** > **Switch Configuration** from the left menu of the Juniper Mist portal. Then select the site that you want to set up, and enter your switch settings.

  If an organization-level switch template was assigned to the site, the site configuration will appear in view-only mode. You can keep the settings from the template or make adjustments. In each section of the page, you can select **Override Configuration Template** and then enter your changes. These changes will apply only to this site, not to the template.

  The following example shows how to override a template and set a site-specific root password.

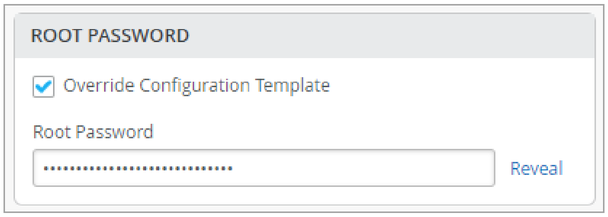
ROOT PASSWORD

☑ Override Configuration Template

Root Password

••••••••••••••••••••••••••••••••

NOTE: The fields that support configuration through site variable have a help text showing the site variable configuration format underneath them. To configure site variables, follow the steps provided in Configure Site Variables. For more information about the switch configuration process and switch templates, see "Configure Switches Using Templates" on page 21.

At both the organization and site levels, the switch settings are grouped into sections as described below.

All Switches

Configure these options in the All Switches section of the **Organization** > **Switch Templates** page and the **Site** > **Switch Configuration** page.

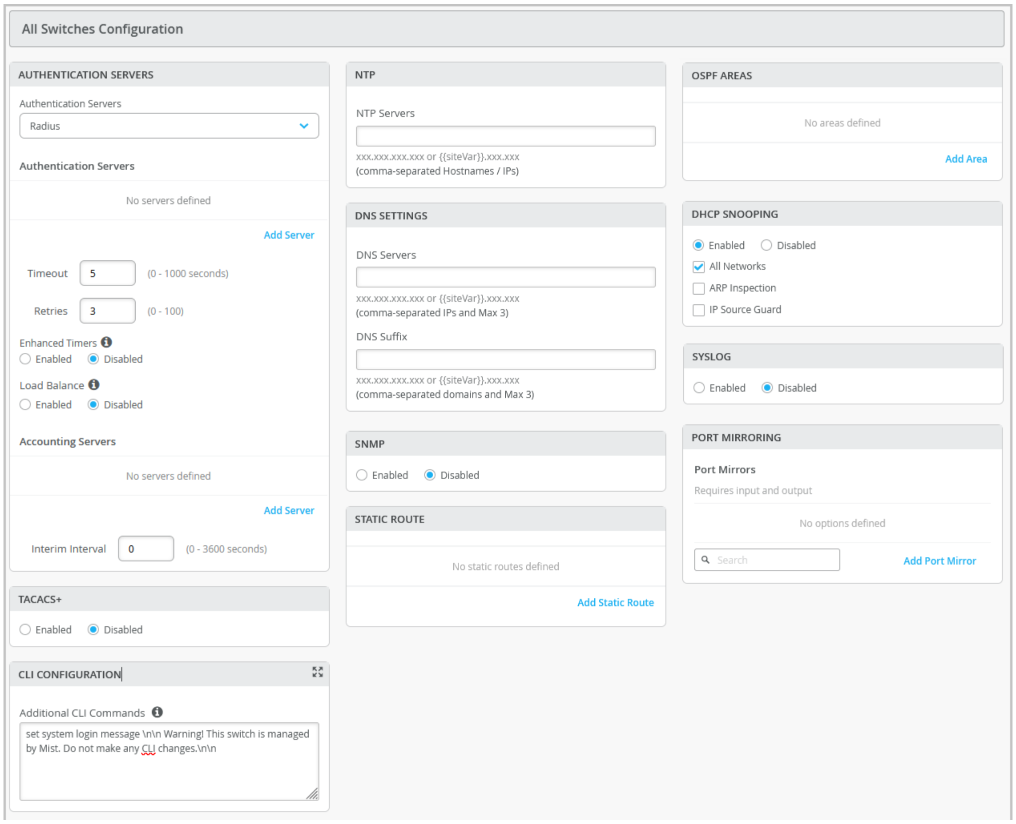

All Switches Configuration

AUTHENTICATION SERVERS

Authentication Servers: Radius

Authentication Servers

No servers defined

Timeout: 5 (0 - 1000 seconds)

Retries: 3 (0 - 100)

Enhanced Timers: ○ Enabled ● Disabled

Load Balance: ○ Enabled ● Disabled

Accounting Servers

No servers defined

Interim Interval: 0 (0 - 3600 seconds)

TACACS+

○ Enabled ● Disabled

CLI CONFIGURATION

Additional CLI Commands

```
set system login message \n\n Warning! This switch is managed by Mist. Do not make any CLI changes.\n\n
```

NTP

NTP Servers

xxx.xxx.xxx.xxx or {{siteVar}}.xxx.xxx
(comma-separated Hostnames / IPs)

DNS SETTINGS

DNS Servers

xxx.xxx.xxx.xxx or {{siteVar}}.xxx.xxx
(comma-separated IPs and Max 3)

DNS Suffix

xxx.xxx.xxx.xxx or {{siteVar}}.xxx.xxx
(comma-separated domains and Max 3)

SNMP

○ Enabled ● Disabled

STATIC ROUTE

No static routes defined

OSPF AREAS

No areas defined

DHCP SNOOPING

● Enabled ○ Disabled

☑ All Networks

☐ ARP Inspection

☐ IP Source Guard

SYSLOG

○ Enabled ● Disabled

PORT MIRRORING

Port Mirrors

Requires input and output

No options defined

Table 2: All Switches Configuration Options

Field/Section	Description
AUTHENTICATION SERVERS	Choose an authentication server for validating usernames and passwords, certificates, or other authentication factors provided by users. • **Mist Auth**—Configure Juniper Mist Access Assurance, a cloud-based authentication service, on your switch. For this option to work, you must use a port with dot1x or MAB authentication. For information, see the Juniper Mist Access Assurance Guide. • **RADIUS**—Select this option to configure a RADIUS authentication server and an accounting server, for enabling dot1x port authentication at the switch level. For the dot1x port authentication to work, you also need to create a port profile that uses dot1x authentication, and you must assign that profile to a port on the switch. The default port numbers are: • port 1812 for the authentication server • port 1813 for the accounting server After selecting an authentication server, configure additional details for the selected server as required. You can configure information that include: • Timeout—Duration in seconds after which the authentication request times out. • Retries—Number of retries allowed. • Enhanced Timers—By default, EX Series switches have a range of 30-60 seconds for various communication timers between the switch and the client device. Enabling this option enhances these timers between 2 and 10 seconds. You can further modify them by changing the authentication server Timeout and Retries.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description

	• Load Balance(Applicable only to RADIUS)—By default EX Series switches use the first RADIUS server. This option randomizes the configuration of the order of servers on a per-switch basis. This ensures load balancing across multiple RADIUS servers.
	• Interim Interval—Specify the frequency (in seconds) at which the authentication server is updated with information about an active user session.
	• Source Address(Applicable only to Mist Auth)—Select a source network. This network should be part of a Layer 3 or IRB interface created with a static IP address.
	• Dynamic Request Port—Specify a change of Authorization (CoA) port.
	NOTE: If you want to set up RADIUS authentication for Switch Management access (for the switch CLI login), you need to include the following CLI commands in the Additional CLI Commands section in the template:
	`set system authentication-order radius`
	`set system radius-server` *`radius-server-IP`* `port 1812`
	`set system radius-server` *`radius-server-IP`* `secret` *`secret-code`*
	`set system radius-server` *`radius-server-IP`* `source-address` *`radius-Source-IP`*
	`set system login user remote class` *`class`*
	For RADIUS or TACACS+ local authentication to the Switch, it is necessary to create a remote user account or a different login class. To use different login classes for different RADIUS-authenticated users, create multiple user templates in the Junos OS configuration by using the following CLI commands in the Additional CLI Commands section:

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description

	`set system login user RO class read-only`
	`set system login user OP class operator`
	`set system login user SU class super-user`
	`set system login user remote full-name "default remote access user template"`
	`set system login user remote class read-only`

TACACS+	Enable TACACS+ for centralized user authentication on network devices. To use TACACS+ authentication on the device, you must configure information about one or more TACACS+ servers on the network. You can also configure TACACS+ accounting on the device to collect statistical data about the users logging in to or out of a LAN and send the data to a TACACS+ accounting server. In addition, you can specify a user role for TACACS+ authenticated users within switch configuration. The following user roles are available: None, Admin, Read, Helpdesk. When the TACACs+ authenticated users do not have a user account configured on the local device, Junos assigns them a user account named 'remote' by default. The port range supported for TACACS+ and accounting servers is 1 to 65535. **NOTE**: For TACACS+ to authenticate into the Switch, a similar login user as defined in the RADIUS section above needs to be created.
NTP	Specify the IP address or hostname of the Network Time Protocol (NTP) server. NTP is used to synchronize the clocks of the switch and other hardware devices on the Internet.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
DNS SETTINGS	Configure the domain name server (DNS) settings. You can configure up to three DNS IP addresses and suffixes in comma separated format.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
SNMP	Configure Simple Network Management Protocol (SNMP) on the switch to support network management and monitoring. You can configure the SNMPv2 or SNMPv3. Here are the SNMP options that you can configure: • **Options under SNMPv2 (V2)** • **General**—Specify the system's name, location, administrative contact information, and a brief description of the managed system. When using SNMPv2, you have the option to specify the source address for SNMP trap packets sent by the device. If you don't specify a source address, the address of the outgoing interface is used by default. • **Client**—Define a list of SNMP clients. You can add multiple client lists. This configuration includes a name for the client list and IP addresses of the clients (in comma separated format). Each client list can have multiple clients. A client is a prefix with /32 mask. • **Trap Group**—Create a named group of hosts to receive the specified trap notifications. At least one trap group must be configured for SNMP traps to be sent. The configuration includes the following fields: • Group Name—Specify a name for the trap group. • Categories—Choose from the following list of categories. You can select multiple values. • authentication

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
	• chassis • configuration • link • remote-operations • routing • services • startup • vrrp-events • Targets—Specify the target IP addresses. You can specify multiple targets. • Version—Specify the version number of SNMP traps. • **Community**—Define an SNMP community. An SNMP community is used to authorize SNMP clients by their source IP address. It also determines the accessibility and permissions (read-only or read-write) for specific MIB objects defined in a view. You can include a client list, authorization information, and a view in the community configuration. • **View**(Applicable to both SNMPv2 and SNMPv3)—Define a MIB view to identify a group of MIB objects. Each object in the view shares a common object identifier (OID) prefix. MIB views allow an agent to have more control over access to specific branches and objects within its MIB tree. A view is made up of a name and a collection of SNMP OIDs, which can be explicitly included or excluded. • **Options under SNMPv3 (V3)**

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
	- **General**—Specify the system's name, location, administrative contact information, and a brief description of the managed system. When using SNMPv2, configure an engine ID, which serves as a unique identifier for SNMPv3 entities. You have an option to use the device MAC address as the engine ID. Using MAC address ensures the engine ID's uniqueness and stability without much manual intervention. - **USM**—Configure the user-based security model (USM) settings. This configuration includes a username, authentication type, and an encryption type. You can configure a local engine or a remote engine for USM. If you select a remote engine, specify an engine identifier in hexadecimal format. This ID is used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host. If you specify the Local Engine option, the engine ID specified on the General tab is considered. If no engine ID is specified, **local mist** is configured as the default value. - **VACM**—Define a view-based access control model (VACM). A VACM lets you set access privileges for a group. You can control access by filtering the MIB objects available for read, write, and notify operations using a predefined view (you must define the required views first from the Views tab). Each view can be associated with a specific security model (v1, v2c, or usm) and security level (authenticated, privacy, or none). You can also apply security settings (you have the option to use already defined USM settings here) to the access group from the Security to Group settings.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
	• **Notify**— Select SNMPv3 management targets for notifications, and specify the notification type. To configure this, assign a name to the notification, choose the targets or tags that should receive the notifications, and indicate whether it should be a trap (unconfirmed) or an inform (confirmed) notification. • **Target**—Configure the message processing and security parameters for sending notifications to a particular management target. You can also specify the target IP address here. • **View**(Applicable to both SNMPv2 and SNMPv3)—Define a MIB view to identify a group of MIB objects. Each object in the view shares a common object identifier (OID) prefix. MIB views allow an agent to have more control over access to specific branches and objects within its MIB tree. A view is made up of a name and a collection of SNMP OIDs, which can be explicitly included or excluded. For more information, see "Configure SNMP on Switches" on page 123.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
STATIC ROUTE	Configure static routes. The switch uses static routes when: - It doesn't have a route with a better (lower) preference value. - It can't determine the route to a destination. - It needs to forward packets that can't be routed. Mist supports IPv4 and IPv6 addresses for static routes. The IPv6 support is available for destination and next hop addresses. Types of static routes supported: - **Subnet**—If you select this option, specify the IP addresses for the destination network and the next hop. - **Network**—If you select this option, specify a VLAN (containing a VLAN ID and a subnet) and the next hop IP address. - **Metric**—The metric value for the static route. This value helps determine the best route among multiple routes to a destination. Range: 0 to 4294967295. - **Preference**—The preference value is used to select routes to destinations in external autonomous systems (ASs) or routing domains. Routes within an AS are selected by the IGP and are based on that protocol’s metric or cost value. Range: 0 to 4294967295. - **Discard**—If you select this check box, packets addressed to this destination are dropped. Discard takes precedence over other parameters. After specifying the details, click the check mark (✓) on the upper right of the **Add Static Route** window to add the configuration to the template.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description

CLI CONFIGURATION	To configure any additional settings that are not available in the template's GUI, you can use **set** CLI commands.
	For instance, you can set up a custom login message to display a warning to users, advising them not to make any CLI changes directly on the switch. Here's an example of how you can do it:
	`set system login message \n\n Warning! This switch is managed by Mist. Do not make any CLI changes.`
	To delete a CLI command that was already added, use the `delete` command, as shown in the following example:
	`delete system login message \n\n Warning! This switch is managed by Mist. Do not make any CLI changes.`
	NOTE: Ensure that you enter the complete CLI command for the configuration to be successful.

OSPF	From this tile, you can: • Define an Open Shortest Path First (OSPF) area. OSPF is a link-state routing protocol used to determine the best path for forwarding IP packets within an IP network. OSPF divides a network into areas to improve scalability and control the flow of routing information. For more information about OSPF areas, see this Junos documentation: Configuring OSPF Areas. • Enable or disable OSPF configuration on the switch (at the switch level). For more information on how to configure OSPF through Mist, refer to "OSPF Configuration for Switches" on page 137.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
VRRP	From this tile you can add a VRRP group by assigning a group number, authentication type, and network(s). For more information, see "Add a VRRP Group to a Configuration" on page 145.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
DHCP SNOOPING	Juniper EX series and QFX series switches provide excellent port security, including DHCP snooping, Address Resolution Protocol (ARP) inspection, and IP Source Guard. You can enable these options for all or selected VLANs on the switch from the Mist portal. DHCP snooping must be enabled for **DHCP** issues to be included in the Wired Successful Connect SLE. DHCP Snooping monitors DHCP messages from untrusted devices connected to the switch. When enabled, DHCP snooping extracts the IP address and lease information from the DHCP packets and stores it in a snooping database. Port security on the EX switches uses this information to verify DHCP requests and block DHCPOFFERs received on untrusted ports (DHCP DISCOVER and DHCP REQUEST are not affected). • **IP Source Guard** works only with single-supplicant 802.1X user authentication mode. It uses the DHCP database to validate source IP addresses and MAC addresses that are received on an untrusted port, and drops those packets that do not have matching entries in the database. • **ARP Inspection** examines the source MAC address in ARP packets received on untrusted ports. It validates the address against the DHCP snooping database, and if the MAC address cannot be found, the packet is dropped. You can use the CLI to check ARP statistics, such as number of invalid ARP packets that it receives on each interface and the sender's IP and MAC addresses, by typing the following commands in the CLI shell: `show dhcp-security arp inspection statistics`, and `show log messages | match DAI` By default, the DHCP protocol considers all trunk ports as trusted and all access ports as untrusted. We recommend that you only connect a DHCP server to the switch using a trunk port, or, if you must use an access port, be sure to explicitly configure that port as trusted in the port profile or DHCP will not work.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
	Note that if you connect a device configured with a static IP address to an untrusted port on the switch, the MAC-IP binding may not exist in the DHCP snooping database; the packets will be dropped. You can use this command `show dhcp-security binding` in a CLI shell to troubleshoot DHCP issues and see what bindings are listed in the DHCP snooping database for the switch. For more information, see DHCP Snooping and Port Security Considerations.
SYSLOG	Configure SYSLOG settings to set up how system log messages are handled. You can configure settings to send the system log messages to files, remote destinations, user terminals, or to the system console. For help with the configuration options, see "Configure the System Log" on page 188.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
PORT MIRRORING	Configure port mirroring. Port mirroring is the ability of a router to send a copy of a packet to an external host address or a packet analyzer for analysis. Mist supports both local and remote port mirroring. In local port mirroring, the source ports and the destination ports (monitor port) are located on the same network switch. In remote port mirroring, the source ports and destination ports are not on the same switch. In this case, the source port forwards the packet copy to the remote destination port through the connection achieved by the ports between the two switches. In the port mirroring configuration, you can specify the following: • **Input**: The source (an interface or network) of the traffic to be monitored. Along with the input, you can specify whether you want Mist to monitor the ingress traffic or the egress traffic for an interface. If you want both ingress and egress traffic to be monitored, add two input entries for the same interface - one with the ingress flag and the other with the egress flag. • **Output**: The destination to which you want to mirror the traffic. You can specify a interface, network, or an IP address (in the case of a remote destination). You cannot specify the same interface or network in both the input and output fields.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
Routing Policy	Configure routing policies for the entire organization (Organization > Switch Templates) or for a site (Site > Switch Configuration). These routing policies will only be pushed to the switch configuration if it is tied to the BGP Routing Protocol. The Routing policies that are already defined inside the BGP tab of a switch will now appear on the Routing Policy tab. The routing policies are tied to protocols such as BGP or OSPF. A routing policy framework is composed of default rules for each routing protocol. These rules determine which routes the protocol places in the routing table and advertises from the routing table. Configuration of a routing policy involves defining terms, which consist of match conditions and actions to apply to matching routes. To configure a routing policy: 1. Click **Add Routing Policy** on the Routing Policy tile. 2. Provide a name to the policy, and then click **Add Terms**. 3. Provide a name to the term and specify other match details such as: • Prefix • AS Path • Protocol • Community—A route attribute used by BGP to administratively group routes with similar properties. • Then—Then action (Accept or Reject) to be applied on the matching routes. • Add Action—Additional actions such as prepend AS path, set community, and set local preference.

Table 2: All Switches Configuration Options *(Continued)*

Field/Section	Description
	4. Click the check mark (✓) on the right of the Add Term title to save the term. You can add multiple terms. 5. Click **Add** to save the routing policy.

Management

Configure these options in the Management section of the **Organization** > **Switch Templates** page and the **Site** > **Switch Configuration** page.

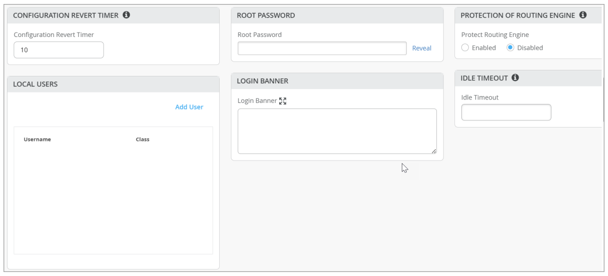

CONFIGURATION REVERT TIMER

Configuration Revert Timer: 10

ROOT PASSWORD

Root Password:

PROTECTION OF ROUTING ENGINE

Protect Routing Engine: ○ Enabled ● Disabled

LOCAL USERS

Username	Class

LOGIN BANNER

Login Banner:

IDLE TIMEOUT

Idle Timeout:

Table 3: Management Configuration Options

Option	Notes
Configuration Revert Timer	This feature helps restore connectivity between a switch and the Mist cloud if a configuration change causes the switch to lose connection. It automatically reverts the changes made by a user and reconnects to the cloud within a specified time duration. By default, this time duration is set to 10 minutes for EX Series switches. You can specify a different time duration. Range: 3 to 30 minutes. In case of a configuration revert event, you can check the switch events page to get specific insight into why the switch configuration was reverted.

Table 3: Management Configuration Options *(Continued)*

Option	Notes
Root Password	A plain-text password for the root-level user (whose username is root).
Protection of Routing Engine	Enable this feature to ensure that the Routing Engine accepts traffic only from trusted systems. This configuration creates a stateless firewall filter that discards all traffic destined for the Routing Engine, except packets from specified trusted sources. Protecting the Routing Engine involves filtering incoming traffic on the router's lo0 interface. Enabling Protection of Routing Engine on Juniper Switches is suggested as the best practice. When Protection of Routing Engine is enabled, Mist by default ensures that the following services (if configured) are allowed to communicate with the switch: BGP, BFD, NTP, DNS, SNMP, TACACS, and RADIUS. If you need additional services that need access to the switch, you can use the Trusted Networks or Services section. If you want to set up access to the switch via ssh, select the ssh option under Trusted Services. If you need to allow switch to respond to pings, select the icmp option under Trusted Services. If you have other segments that you would like to reach the switch from, you can add them under Trusted Networks or Trusted IP/Port/Protocol. For more information, refer to Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources and Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods.
Local Users	Create a local user account on the switch for device management purposes. To create a user account, click **Add User** and then define a username, login class (Operator, Read-only, Super User, or Unauthorized), and a password.

Table 3: Management Configuration Options *(Continued)*

Option	Notes
Idle Timeout	The maximum number of minutes that a remote shell session can be idle. When this limit is reached, users are logged out. (Valid Range: 1-60).
Login Banner	Enter text that you want users to see when they log in to the switch. Example: “Warning! This switch is managed by Juniper Mist. Do not make any CLI changes.” You can enter up to 2048 characters.
DHCP Option 81 (For Dynamic DNS)	Enable switches with DHCP option 81 support. When this option is enabled on a switch, the clients connected to that switch can send their fully qualified domain name (FQDN) to the DHCP server while requesting an IP address. This allows the DHCP server to update DNS records accordingly. You can enable the DHCP option 81 at the site level (**Site** > **Switch Configuration**) and device level (**Switches** > ***Switch Name***) as well.

Shared Elements

Configure these options in the Shared Elements section of the **Organization** > **Switch Templates** page and the **Site** > **Switch Configuration** page.

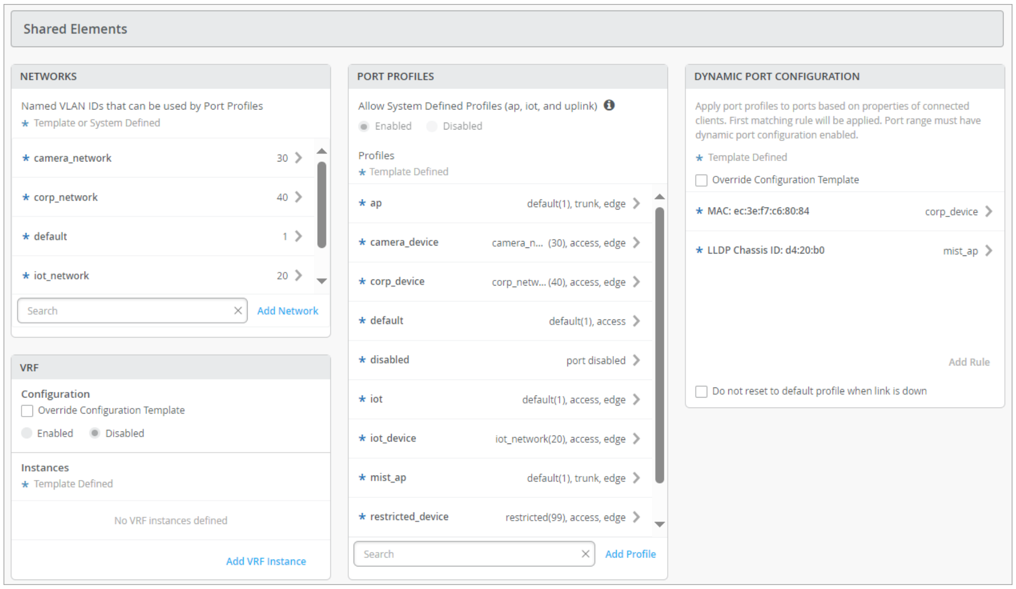

Shared Elements

NETWORKS

Named VLAN IDs that can be used by Port Profiles

* Template or System Defined

* camera_network	30
* corp_network	40
* default	1
* iot_network	20

VRF

Configuration

□ Override Configuration Template

○ Enabled ● Disabled

Instances

* Template Defined

No VRF instances defined

PORT PROFILES

Allow System Defined Profiles (ap, iot, and uplink)

● Enabled ○ Disabled

Profiles

* Template Defined

* ap	default(1), trunk, edge
* camera_device	camera_n... (30), access, edge
* corp_device	corp_netw... (40), access, edge
* default	default(1), access
* disabled	port disabled
* iot	default(1), access, edge
* iot_device	iot_network(20), access, edge
* mist_ap	default(1), trunk, edge
* restricted_device	restricted(99), access, edge

DYNAMIC PORT CONFIGURATION

Apply port profiles to ports based on properties of connected clients. First matching rule will be applied. Port range must have dynamic port configuration enabled.

* Template Defined

□ Override Configuration Template

* MAC: ec:3e:f7:c6:80:84	corp_device
* LLDP Chassis ID: d4:20:b0	mist_ap

□ Do not reset to default profile when link is down

Table 4: Shared Elements Configuration Options

Option	Notes
Networks	Add or update VLANs, which you can then use in your port profiles. For each VLAN, enter the name, VLAN ID, and subnet. You can specify IPv4 or IPv6 address for the subnet. See the on-screen information for more tips. On this tile, you have an option to hide the networks that are not used in a user-defined port profiles or L3 sub-interfaces. This feature helps you quickly identify those networks that are in use and those that are not in use.
Port Profiles	Add or update port profiles. For help with the profile options, see the on-screen tips and "Shared Elements—Port Profiles" on page 52. On this tile, you have an option to hide the port profiles that are not used in any static or dynamic port configurations defined by users. This feature helps you quickly identify those port profiles that are in use and those that are not in use.

Table 4: Shared Elements Configuration Options *(Continued)*

Option	Notes
Dynamic Port Configuration	Dynamic port profiling uses a set of device properties of the connected client device to automatically associate pre-configured port and network settings to the interface. Dynamic port profile configuration involves the following two steps at a high level: • Configure a dynamic port profile rules (described here). • Specify the ports that you want to function as dynamic ports. You can do this by selecting the **Enable Dynamic Configuration** check box on the Port Config tab in the Select Switches section of the switch template or in the Port Configuration section of the switch details page. For more information, see the Enable Dynamic Configuration row in Table 7 on page 62. You can configure a dynamic port profile rules using the following parameters: • LLDP System Name • LLDP Description • LLDP Chassis ID • Radius Username • Radius Filter-ID • MAC (Ethernet mac-address) In this example, the port profile specified in the **Apply Configuration Profile** field will be assigned to a switch port enabled with dynamic configuration when it is connected to any devices with an LLDP system name that matches the parameters configured.

Table 4: Shared Elements Configuration Options *(Continued)*

Option	Notes
	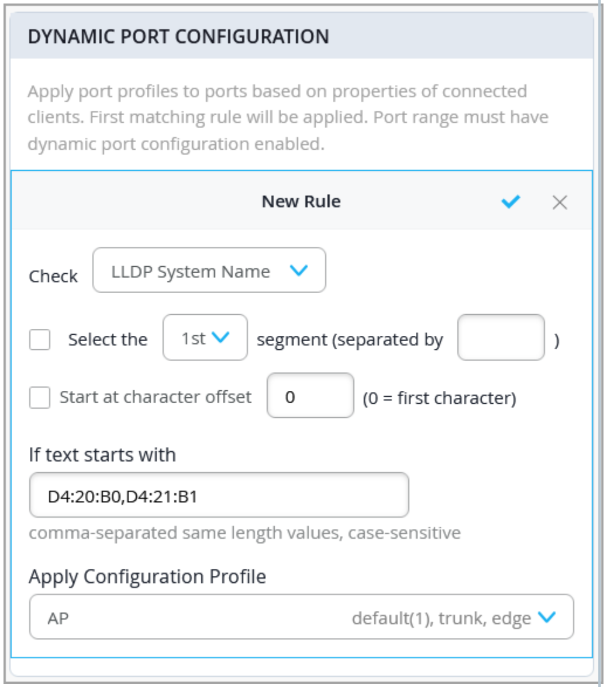 DYNAMIC PORT CONFIGURATION Apply port profiles to ports based on properties of connected clients. First matching rule will be applied. Port range must have dynamic port configuration enabled. New Rule ✓ × Check: LLDP System Name □ Select the 1st segment (separated by) □ Start at character offset 0 (0 = first character) If text starts with: D4:20:B0,D4:21:B1 comma-separated same length values, case-sensitive Apply Configuration Profile: AP default(1), trunk, edge **NOTE**: • If you use multiple values in the **If text starts with** field in a DPC rule, separate them with commas and ensure that they all have the same length. If any value differs in length, you must create a separate rule for it. • Prefer LLDP-based matching over MAC-based matching when the device supports LLDP. • Do not use MAC-based matching on ports enabled with 802.1X authentication. • Avoid using `Filter-ID` attributes. When 802.1X is enabled on the ports, VLAN assignment should be handled via RADIUS without relying on `Filter-ID`. For more information, refer to "Configure Dynamic Port Profile Assignment" on page 74.

Table 4: Shared Elements Configuration Options *(Continued)*

Option	Notes
VRF	With VRF, you can divide an EX Series switch into multiple virtual routing instances, effectively isolating the traffic within the network. You can define a name for the VRF, specify the networks associated with it, and include any additional routes needed. You can specify IPv4 or IPv6 addresses for the additional route. **NOTE**: • You can't assign the default network (VLAN ID = 1) to VRF. • Mist recommends using VRFs in network segments where traffic isolation and overlapping IP address spaces are required.

Shared Elements—Port Profiles

In the "Shared Elements" on page 48 section, you can configure port profiles. These options appear when you click **Add Profile** or when you click a profile to edit.

NOTE:

- For general information about profiles, see "Port Profiles" on page 71.
- If you're working at the site level, you might see asterisks (*) next to the port profile names. These port profiles were created in the switch template. If you click them, you'll see the settings in view-only mode. To make site-specific changes (affecting only this site and not the switch template itself), select **Override Template Defined Profile** and then edit the settings.

Table 5: Port Profile Configuration Options

Option	Notes
Name, Port Enabled, and Description	Basic settings to identify and enable the port.
Mode	• Trunk—Trunk interfaces typically connect to other switches, APs, and routers on the LAN. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Specify the Port Network, VoIP Network (if applicable), and Trunk Networks. • Access—Default mode. Access interfaces typically connect to network devices, such as PCs, printers, IP phones, and IP cameras. In this mode, the interface can be in a single VLAN only.
Port Network (Untagged/Native VLAN)	Specify the Port Network or native VLAN.
VoIP Network	Specify the VoIP Network (if applicable).
Trunk Networks	Specify a trunk network if you have chosen the mode Trunk.

Table 5: Port Profile Configuration Options *(Continued)*

Option	Notes
Use dot1x authentication	Select this option to enable IEEE 802.1X authentication for Port-Based Network Access Control. 802.1X authentication is supported on interfaces that are members of private VLANs (PVLANs). The following options are available if you enable dot1x authentication on a port: • **Allow Multiple Supplicants**—Select this option to allow multiple end devices to connect to the port. Each device is authenticated individually. • **Dynamic VLAN**—Specify dynamic VLANs that will be returned by the RADIUS server attribute 'tunnel-private-group-ID' or 'Egress-VLAN-Name'. This configuration enables a port to perform dynamic VLAN assignment. • **MAC authentication**—Select this option to enable MAC authentication for the port. When this option is selected, you can also specify an **Authentication Protocol**. If you specify a protocol, it must be used by supplicants to provide authentication credentials. • **Use Guest Network**—Select this option to use a guest network for authentication. Then select a **Guest Network** from the drop-down list. • **Bypass authentication when server is down**—If you select this option, clients can join the network without authentication if the server is down. • **Reauthentication interval**—In a switch port profile that uses dot1x authentication, you can configure a timer that controls how often a client reauthenticates itself with the RADIUS server. The recommended value is 6 to 12 hours (21600 to 43200 seconds). The default value is 65000 seconds. • **Server Reject Network**—Select this option to connect users to a specified VLAN (such as the guest network) in the event that the authentication server rejects the user-authentication attempt. You can configure this option at the switch-level, site template-level, or organization template-level. • **Server Fail Network**—Select this option to connect users to a specified VLAN (such as the guest network) in the event that the authentication server cannot be reached or fails to respond. You can configure this option at the switch-level, site template-level, or organization template-level. You need to also do the following for dot1x authentication to work:

Table 5: Port Profile Configuration Options *(Continued)*

Option	Notes
	• Configure a RADIUS server for dot1x authentication from the Authentication Servers tile in the All Switches Configuration section of the template. • Assign a dot1x port profile to a switch port for the RADIUS configuration to be pushed to the switch. You can do this from the Port Config tab in the Select Switches Configuration section of the template. Mouse-over the port to see the RADIUS-assigned VLAN field. Ports with dot1x enabled are assigned a new VLAN by the RADIUS server when the 802.1x authentication is successful. This view is especially useful when checking whether a given VLAN on a port has changed following dot1x authentication.

Table 5: Port Profile Configuration Options *(Continued)*

Option	Notes
	Figure 1: Radius Assigned VLAN on a Dot1x Port 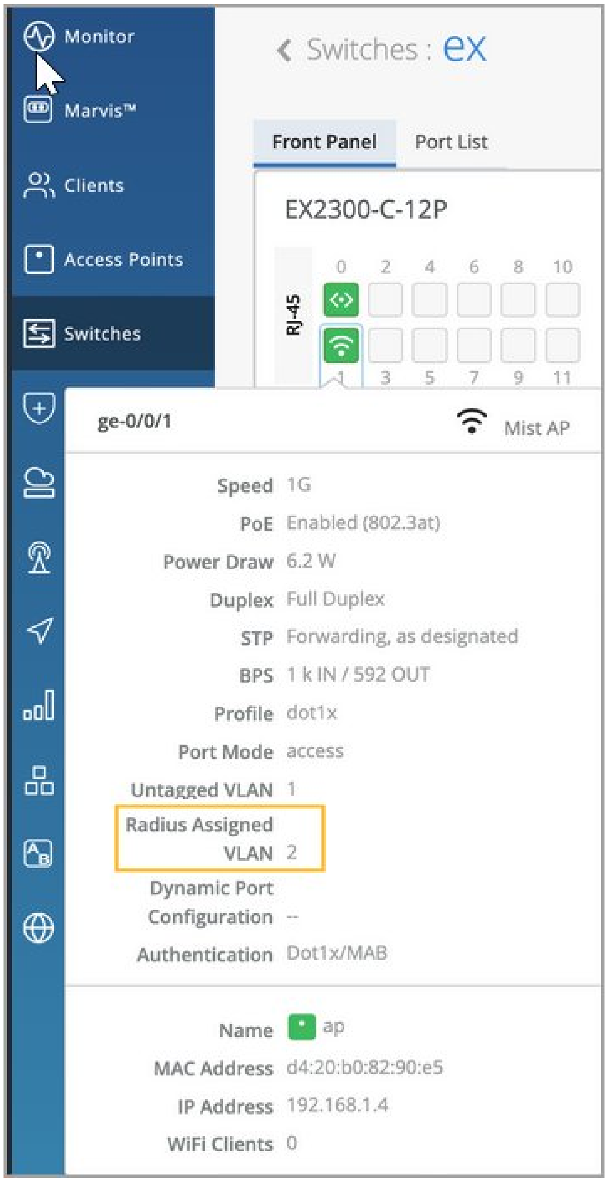 Monitor Marvis™ Clients Access Points Switches Switches : ex Front Panel Port List RJ-45 0 2 4 6 8 10 1 3 5 7 9 11 EX2300-C-12P ge-0/0/1 Mist AP Speed 1G PoE Enabled (802.3at) Power Draw 6.2 W Duplex Full Duplex STP Forwarding, as designated BPS 1 k IN / 592 OUT Profile dot1x Port Mode access Untagged VLAN 1 Radius Assigned VLAN 2 Dynamic Port Configuration -- Authentication Dot1x/MAB Name ap MAC Address d4:20:b0:82:90:e5 IP Address 192.168.1.4 WiFi Clients 0
Speed	Keep the default setting, Auto, or select a speed
Duplex	Keep the default setting, Auto, or select a Half or Full.

Table 5: Port Profile Configuration Options *(Continued)*

Option	Notes
MAC Limit	Configure the maximum number of MAC addresses that can be dynamically learned by an interface. When the interface exceeds the configured MAC limit, it drops the frames. A MAC limit also results in a log entry. The configured value remains active until it is replaced or cleared, and persists through device reboot. The default value: 0 Supported range: 0 through 16383
PoE	Enable the port to support power over Ethernet (PoE).
Per VLAN STP	Configure a switch with VLAN Spanning Tree Protocol (VSTP) or per-VLAN Spanning Tree. VSTP helps in preventing loops in Layer 2 networks on a per-VLAN basis. One Spanning Tree per VLAN enables fine grain load balancing. Mist recommends enabling this feature for other vendor's devices (for example, Cisco) that operate per-VLAN Spanning Tree by default. This setting is available at the site and switch level as well.
STP Edge	Configure the port as a Spanning Tree Protocol (STP) edge port, if you want to enable Bridge Protocol Data Unit (BPDU) guard on a port. STP Edge is enabled on ports to which clients that do not participate in STP are connected. This setting ensures that the port is treated as an edge port and guards against the reception of BPDUs. If you plug a non-edge device into a port configured with STP Edge, the port is disabled. In addition, the Switch Insights page generates a Port BPDU Blocked event. The Front Panel on the "Switch Details" on page 294 will also display a BPDU Error for this port. You can clear the port of the BPDU error by selecting the port on the Front Panel and then clicking **Clear BPDU Errors**. You should not enable STP Edge on the Uplink port. You can also configure STP Edge at the switch level, from the Port Profile section on the switch details page.
STP Point-to-Point	This configuration changes the interface mode to point-to-point. Point-to-point links are dedicated links between two network nodes, or switches, that connect one port to another.
STP No Root Port	This configuration prevents the interface from becoming a root port.

Table 5: Port Profile Configuration Options *(Continued)*

Option	Notes
Block STP BPDUs	Typically enabled on edge or access ports where BPDUs are not expected. When this option is enabled, the port is immediately shut down if a BPDU is received, helping to prevent potential loops or misconfigurations. If STP Edge is enabled, Block STP BPDUs is automatically disabled. However, BPDU Liveliness Check can still be configured. If Block STP BPDUs is enabled, both STP Edge and BPDU Liveliness Check are automatically disabled.
STP BPDU Liveliness Check	Typically enabled on uplink or trunk ports where BPDUs are expected. This feature monitors BPDU reception and blocks the port and raises an alarm if no BPDUs are received within 20 seconds, helping detect failures or misconfigurations quickly.
QoS	Enable Quality of Service (QoS) for the port to prioritize latency-sensitive traffic, such as voice, over other traffic on a port. **NOTE**: For optimal results, it's important to enable Quality of Service (QoS) for both the downstream (incoming) and upstream (outgoing) traffic. This ensures that the network can effectively prioritize and manage traffic in both directions, leading to improved performance and better overall quality of service. You have the option to override the QoS configuration on the WLAN settings page (**Site** > **WLANs** > ***WLAN name***). To override the QoS configuration, select the **Override QoS** check box and choose a wireless access class. The downstream traffic (AP > client) gets marked with the override access class value specified. The override configuration doesn't support upstream traffic (client > AP). See also: "QoS Configuration" on page 110.
Storm Control	Enable storm control to monitor traffic levels and automatically drop broadcast, multicast, and unknown unicast packets when the traffic exceeds a traffic level (specified in percentage). This specified traffic level is known as the storm control level. This feature actively prevents packet proliferation and maintains the performance of the LAN. When you enable Storm Control, you can also choose to exclude broadcast, multicast, and unknown unicast packets from monitoring. You can also configure a switch to automatically shut down a port when traffic exceeds the user-defined storm control threshold, by selecting the **Shutdown Port** check box under Action on Threshold. For more information, see Understanding Storm Control.

Table 5: Port Profile Configuration Options *(Continued)*

Option	Notes
Persistent (Sticky) MAC Learning	Enable Persistent (Sticky) MAC to retain MAC addresses for trusted workstations and servers learned by the interface, even after a device restart. You can configure Sticky MAC for static wired clients. Sticky MAC is not intended for use on Juniper Mist AP interfaces, nor is it supported for trunk ports or those configured with 802.1X authentication. Used in conjunction with MAC Limits (explained above), Sticky MAC protects against Layer 2 denial-of-service (DoS) attacks, overflow attacks on the Ethernet switching table, and DHCP starvation attacks, while still allowing the interface to dynamically learn MAC addresses. In the Mist portal, the Insights page reports these events as *MAC Limit Exceeded* . You configure both Sticky MAC and MAC limits as part of the Port Profile for the switch. The general procedure is demonstrated in this video: **Video:** Port Profiles You must explicitly enable the **Persistent (Sticky) MAC Learning** option, located at the bottom of the Port Profile configuration block, to include Sticky MAC as part of the Port Profile that you associate with the interface. For MAC limits, the default value is 0 (unlimited, that is, disabled) but you can enable it by setting a value of up to 16383 unique MAC addresses allowed. To see in the Mist portal what value has been set for the MAC Limit or the MAC Count, select a switch from the Switches page and hover your mouse over a switch port. You can see which (port) Profile is applied to the interface, and by extension, know its Sticky MAC status.

Table 5: Port Profile Configuration Options *(Continued)*

<table>
<tr><th>Option</th><th>Notes</th></tr>
<tr><td></td><td>

Figure 2: Port Details Showing Sticky MAC

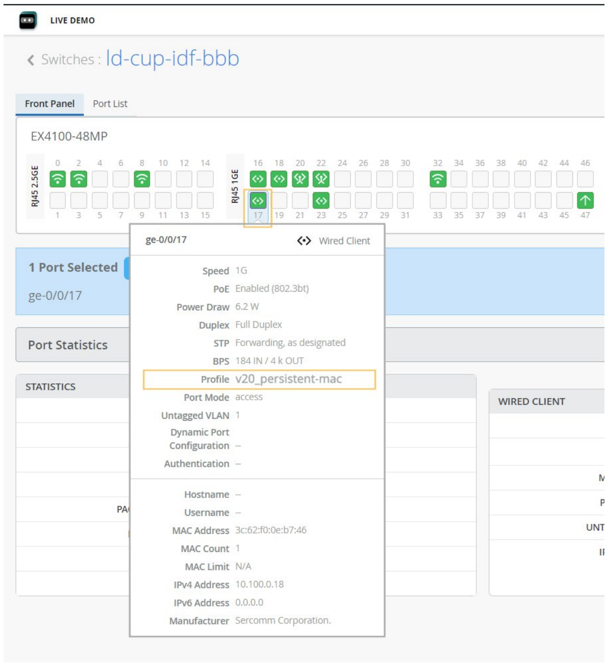

ge-0/0/17 — Wired Client

Speed 1G
PoE Enabled (802.3bt)
Power Draw 6.2 W
Duplex Full Duplex
STP Forwarding, as designated
BPS 184 IN / 4 k OUT
Profile v20_persistent-mac
Port Mode access
Untagged VLAN 1
Dynamic Port Configuration --
Authentication --

Hostname --
Username --
MAC Address 3c:62:f0:0e:b7:46
MAC Count 1
MAC Limit N/A
IPv4 Address 10.100.0.18
IPv6 Address 0.0.0.0
Manufacturer Sercomm Corporation.

The configured MAC limit and number of MACs learned will appear after a few minutes, as dynamic learning on the interface progresses. In the Mist dashboard, only the maximum MAC address count is shown. However, you can see every MAC address a given interface has learned by opening a **Remote Shell** to the switch and running the following Junos CLI commands:

```
show ethernet-switching table persistent-learning
show ethernet-switching table persistent-learning interface
```

MAC count is a persistent value that remains until the MAC address is cleared (or until it is disabled in the Port Profile and then that configuration is pushed to the switch).

To clear the MAC addresses on a given interface from the Mist dashboard, you need to be logged as Network Administrator or Super User. Then just select the port you want from the switch front panel (as shown in Figure 1) and click the **Clear MAC [Dynamic/ Persistent]** button that appears.

On the Switch Insights page, the event shows up as a *MAC Limit Reset* event.

</td></tr>
</table>

Table 5: Port Profile Configuration Options *(Continued)*

Option	Notes
	For more information on the front panel, see "Switch Details" on page 294.

Select Switches Configuration

Create rules to apply configuration settings based on the name, role, or model of the switch.

Click a rule to edit it, or click **Add Rule**. Then complete each tabbed page. As you enter settings, click the checkmark at the top right to save your changes. You can also create a switch rule entry by cloning an existing rule. To do that, you just need to click the clone button and name the new rule.

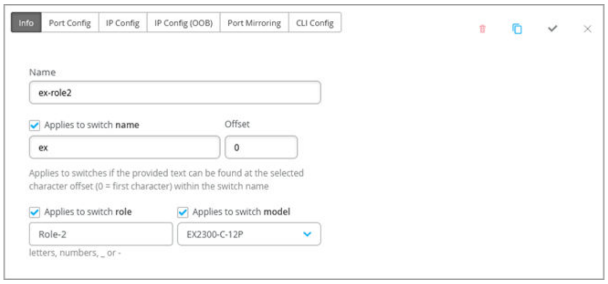

Name: ex-role2

☑ Applies to switch **name**: ex
Offset: 0

Applies to switches if the provided text can be found at the selected character offset (0 = first character) within the switch name

☑ Applies to switch **role**: Role-2
letters, numbers, _ or -

☑ Applies to switch **model**: EX2300-C-12P

The various tabs are described in separate tables below.

Table 6: Select Switches—Info Tab

Option	Notes
Name	Enter a name to identify this rule.
Applies to switch name	Enable this option if you want this rule to apply to all switches that match the specified name. Then enter the text and the number of offset characters. For example, if you enter *abc* with an offset of 0, the rule applies to switches whose names start with *abc*. If the offset is 5, the rule ignores the first 5 characters of the switch name.

Table 6: Select Switches—Info Tab *(Continued)*

Option	Notes
Applies to switch role	Enable this option if you want this rule to apply to all switches that have the same role. Enter the role by using lowercase letters, numbers, underscores (_), or dashes (-).
Applies to switch model	Enable this option if you want this rule to apply to all switches that have the same model. Then select the model.

Table 7: Select Switches—Port Config Tab

Option	Notes
Configuration List	Click **Add Port Configuration**, or select a port configuration to edit. 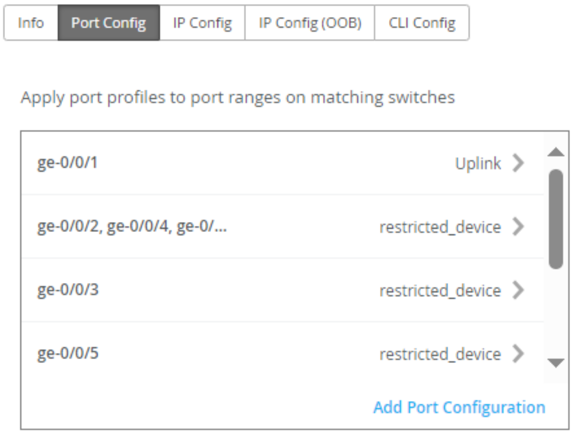 Info Port Config IP Config IP Config (OOB) CLI Config Apply port profiles to port ranges on matching switches ge-0/0/1 Uplink ge-0/0/2, ge-0/0/4, ge-0/... restricted_device ge-0/0/3 restricted_device ge-0/0/5 restricted_device
Port IDs	Enter the port(s) to configure.
Configuration Profile	Select the configuration profile to apply to the specified ports. **NOTE**: If you want to configure switch ports with Q-in-Q tunneling, choose Q-in-Q from this drop-down list. For more information, refer to "Configure Q-in-Q Tunneling on a Switch Port" on page 148.

Table 7: Select Switches—Port Config Tab *(Continued)*

Option	Notes
Port Network (S-VLAN)	Specify a service VLAN (S-VLAN) if the port is using Q-in-Q tunneling. S-VLAN is an external, additional VLAN tag used to extend Layer 2 Ethernet connections between customer sites. This is especially useful when customers have overlapping VLAN IDs.
Speed	**NOTE**: Applicable only if you have selected Q-in-Q as the configuration profile. Keep the default setting, Auto, or select a speed.
Duplex	**NOTE**: Applicable only if you have selected Q-in-Q as the configuration profile. Keep the default setting, Auto, or select a Half or Full.
PoE	**NOTE**: Applicable only if you have selected Q-in-Q as the configuration profile. Enable the port to support power over Ethernet (PoE).
MTU	**NOTE**: Applicable only if you have selected Q-in-Q as the configuration profile. Specify the media maximum transmission unit (MTU) for the port. Default: 1514. Range: 256 - 9216. The media maximum transmission unit (MTU) for an interface is the *largest data unit that can be forwarded through that interface* without fragmentation.

Table 7: Select Switches—Port Config Tab *(Continued)*

Option	Notes
Storm Control	**NOTE**: Applicable only if you have selected Q-in-Q as the configuration profile. Enable storm control to monitor traffic levels and automatically drop broadcast, multicast, and unknown unicast packets when the traffic exceeds a traffic level (specified in percentage). This specified traffic level is known as the storm control level. This feature actively prevents packet proliferation and maintains the performance of the LAN. When you enable Storm Control, you can also choose to exclude broadcast, multicast, and unknown unicast packets from monitoring. You can also configure a switch to automatically shut down a port when traffic exceeds the user-defined storm control threshold, by selecting the **Shutdown Port** check box under Action on Threshold. For more information, see Understanding Storm Control.
Description	Provide a description for the port.

Table 7: Select Switches—Port Config Tab *(Continued)*

Option	Notes
Enable Dynamic Configuration	(This setting is not applicable if you have selected Q-in-Q as the configuration profile.) **NOTE**: Ensure that you have created a restricted VLAN and network profile that can be assigned to unknown devices that are connected to a switch port enabled with dynamic port configuration but do not match the dynamic port assignment rules. This setting enables a switch port to work as a dynamic port, which uses the dynamic port assignment rules (described in the **Dynamic Port Configuration** row in Table 4 on page 49). When a device is connected to a switch port enabled with dynamic port configuration, a port profile is dynamically assigned to it based on attributes of the connected device. If the device matches the attributes, Mist assigns a matching dynamic profile to the device. But if the device doesn't match the attributes, it is assigned a specified VLAN, ideally a restricted VLAN (port profile). In the following example, the port is enabled with dynamic port allocation and is assigned with a restricted VLAN. In this case, if the connected device doesn't match the dynamic profiling attributes, it will be placed into a restricted VLAN such as a non-routable VLAN or a guest VLAN. Interfaces enabled with Port Aggregation don't support dynamic port configuration.

Table 7: Select Switches—Port Config Tab *(Continued)*

Option	Notes
	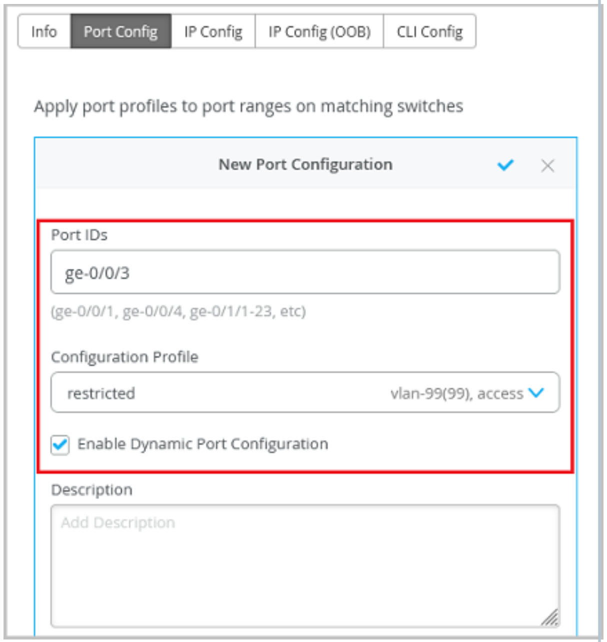 Info Port Config IP Config IP Config (OOB) CLI Config Apply port profiles to port ranges on matching switches New Port Configuration Port IDs: ge-0/0/3 (ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc) Configuration Profile: restricted vlan-99(99), access ☑ Enable Dynamic Port Configuration Description: Add Description It takes a couple of minutes for a port profile to be applied a port after a client is recognized, and a couple of minutes after that for the port profile assignment status to appear on the Mist portal. In case of switch reboots or a mass link up or down event affecting all ports on a switch, it takes approximately 20 minutes for all the ports to be assigned to the right profile (assuming that dynamic port configuration is enabled on all the ports). Dynamic port configuration on a switch is meant for establishing connection to IoT devices, APs, and user port endpoints. Do not use it for creating connection between switches, switches and routers, and switches and firewalls. Also, you should not enable Dynamic Port Configuration on the uplink port. **NOTE**: • Ensure that the default or restricted VLAN used in dynamic port configuration does not have an active DHCP server running. Otherwise, you

Table 7: Select Switches—Port Config Tab *(Continued)*

Option	Notes
	might encounter stale IP address issue on certain legacy devices. - A switch with port-based network access control (NAC) authentication does not require dynamic port configuration as VLAN assignments are handled by the RADIUS server. Also, we do not recommend using dynamic port profiles when RADIUS server with MAC Authentication Bypass (MAB) is used. For more information, refer to "Configure Dynamic Port Profile Assignment" on page 74.
Up/Down Port Alerts	When you enable this feature, Juniper Mist monitors transitions between up and down states on these ports. If you enable this feature, also enable Critical Switch Port Up/Down on the Monitor > Alerts > Alerts Configuration page.

Table 7: Select Switches—Port Config Tab ***(Continued)***

Option	Notes
Port Aggregation	**NOTE**: Not applicable if you have selected Q-in-Q as the configuration profile. When you enable this feature, the Ethernet interfaces specified are grouped to form a single link layer interface. This interface is also known as a link aggregation group (LAG) or bundle. The number of interfaces that you can group into a LAG and the total number of LAGs that a switch supports vary depending on switch model. You can use LAG with or without LACP enabled. If the device on the other end doesn't support LACP, you can disable LACP here. You can also specify the following: • The LACP force-up state for the switch. This configuration sets the state of the interface as up when the peer has limited LACP capability. • An LACP packet transmission interval. If you configure the LACP Periodic Slow option on an AE interface, the LACP packets are transmitted every 30 seconds. By default, the interval is set to fast in which the packets are transmitted every second. • An AE index. Ensure that the AE index does not overlap across different ports between the device, site or template, and campus fabric configuration. For more information on how to configure link aggregation group (LAG) with Wired Assurance, watch the following video: **Video:**
Allow switch port operator to modify port profile	When you enable this feature, users with the Switch Port Operator admin role can view and manage this configuration.

Table 8: Select Switches Configuration—IP Config Tab

Option	Notes
Network (VLAN) List	Select a network for in-band management traffic. Or click Add Network and complete the New Network fields as described in the remaining rows of this table. 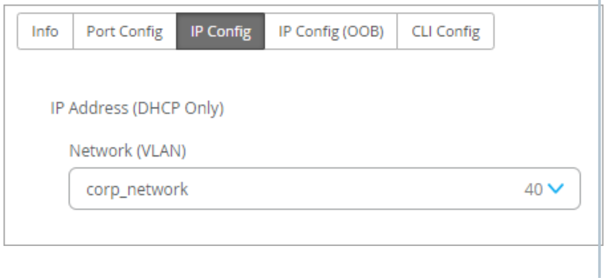 IP Address (DHCP Only) Network (VLAN) corp_network 40
Name	Enter a name to identify this network.
VLAN ID	Enter the VLAN ID from 1-4094, or enter a site variable to dynamically enter an ID.
Subnet	Enter the subnet or site variable.

Select Switches—IP Config (OOB) Tab

Enable or disable **Dedicated Management VRF** (out of band). For all standalone devices or Virtual Chassis running Junos version 21.4 or later, this feature confines the management interface to non-default virtual routing and forwarding (VRF) instances. Management traffic no longer has to share a routing table with other control traffic or protocol traffic.

Select Switches—Port Mirroring Tab

This tab displays the list of port mirroring configurations already added. Click an entry to edit it. Or click **Add Port Mirror** to enable port mirroring. This feature allows you to dynamically apply port mirroring on switches based on the parameters such as the switch role, switch name, and switch model as specified in the rules. This feature is typically used for monitoring and troubleshooting. When port mirroring is enabled, the switch sends a copy of the network packet from the mirrored ports to the monitor port.

Mist supports both local and remote port mirroring. In local port mirroring, the source ports and the destination ports (monitor port) are located on the same network switch. In remote port mirroring, the source ports and destination ports are not on the same switch. In this case, the source port forwards the packet copy to the remote destination port through the connection achieved by the ports between the two switches.

The configuration options include the following:

- **Input**—The source (an interface or network) of the traffic to be monitored. Along with the input, you can specify whether you want Mist to monitor the ingress traffic or the egress traffic for an interface. If you want both ingress and egress traffic to be monitored, add two input entries for the same interface - one with the ingress flag and the other with the egress flag.

- **Output**—The destination to which you want to mirror the traffic. You can specify an interface, network, or an IP address (in case of a remote destination). You cannot specify the same interface or network in both the input and output fields.

The rules under Select Switches Configuration take precedence over the global Port Mirroring configuration. Also, if the global port mirroring is configured, it is displayed as the default rule in the Select Switches configuration section and is displayed as read-only. You can edit it at the global level.

Select Switches—CLI Config Tab

Enter additional CLI commands, as needed.

Switch Policy Labels, GBP Tags, and Switch Policies

Use this section to create Access Control Lists (ACLs) (also known as firewall filters) and Group-Based Policies (GBP).

- Source/Destination labels—Create labels to identify the source/destination IP addresses for Access Control List (ACL) policies (RADIUS-based firewall filters). For more information, see "Firewall Filters" on page 195.

- GBP tags—(For Campus Fabric IP-Clos deployments) Create tags for Group-Based Policies (GBP), which leverage VXLAN technology. GBP simplifies configuration and provides endpoint access control across your campus. For more information, see "Group-Based Policies" on page 286.

Port Profiles

IN THIS SECTION

Port profiles provide a convenient way to manually or automatically provision switch interfaces. Mist supports the following types of port profiles:

- System-defined—System-defined Port Profiles are port profiles that are built into the Mist portal and are available for you to use if you do not want to configure your own port profiles. These are preconfigured for you, so there is no configuration required in order for you to be able to use them. The system-defined port profiles provided by Mist include the following: ap, iot, uplink, default, and disabled.
- User-defined—You can create these custom port profiles based on your use case if the system-defined profiles do not meet your requirements.

At a high level, port profile configuration involves the following two steps: defining a port profile (or using an existing system-defined profile) and assigning it to switch ports.

Define Port Profiles

You can either use an existing profile or define one that meets your requirements. Mist supports creating port profiles at the switch template level or at the individual switch level.

To define a port profile:

1. Go to **Organization** > **Switch Templates** and the select the appropriate switch template.

   If you wish to define a port profile at the switch level, navigate to the switch details page (**Switches** > *switch-name*).

> **NOTE**: Before creating a port profile, ensure that you have created the required networks from the **Networks** tile. In a switch template, you can find this tile in the Shared Elements section. On a switch details page, you can find this tile in the Networks & Port Profiles section.

2. Navigate to the **Port Profile** tile.

   In a switch template, you can find this tile in the Shared Elements section. On a switch details page (for switch-level action), you can find this tile in the Networks & Port Profiles section.

3. On the Port Profile tile, click **Add Profile**.

   The **New Port Profile** window appears.

4. Specify the profile settings as described in Table 5 on page 53.

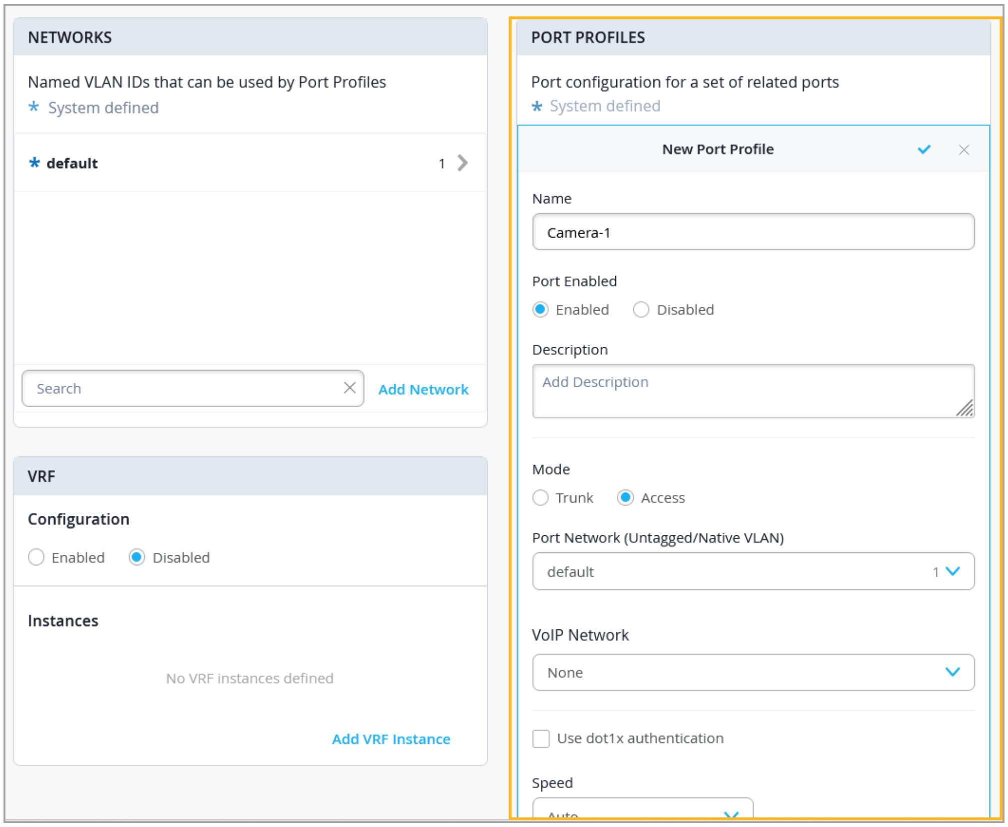

NETWORKS

Named VLAN IDs that can be used by Port Profiles
* System defined

* default 1 >

Search ✕ Add Network

VRF

Configuration
○ Enabled ● Disabled

Instances

No VRF instances defined

Add VRF Instance

PORT PROFILES

Port configuration for a set of related ports
* System defined

New Port Profile ✓ ✕

Name
Camera-1

Port Enabled
● Enabled ○ Disabled

Description
Add Description

Mode
○ Trunk ● Access

Port Network (Untagged/Native VLAN)
default 1

VoIP Network
None

□ Use dot1x authentication

Speed
Auto

5. After specifying the settings, click the check mark (✓) on the upper right of the **New Port Profile** window.

   The newly created port profile appears on the list of profiles within the Port Profile tile.

6. Click **Save** on the upper right of the template or switch details page to confirm and save the changes.

Assign Port Profiles to Switch Ports

IN THIS SECTION

Based on how a profile is assigned to a switch port, Mist supports the following types of port profiles:

- Static port profiles—A static port profile is the profile that is manually assigned to a specific switch port. These profiles are used for static provisioning of switch ports.

- Dynamic port profiles—Dynamic port profiles help the switch port detect the device connected to it by using the port assignment rules configured and assign a matching profile to the port dynamically. Dynamic port profiles are used for autoprovisioning of switch ports (colorless ports).

Assign Port Profiles Manually

After you define a port profile, you can assign it manually to a specific switch port from the Port Config tab in the Select Switches section of the switch template, or from the Port Configuration section on the switch details page.

To manually assign a port profile to a port via a switch template:

1. Go to **Organization** > **Wired** > **Switch Templates** and then click a template to open it.
2. Navigate to the Select Switches Configuration section.
3. Click an existing rule to open it.

   If you want to create a new rule with port assignment settings, click **Add Rule** on the upper right of the section.
4. Go to the **Port Config** tab of the rule.
5. Click **Add Port Configuration**.

   The New Port Configuration window appears.
6. Specify the ports in the **Port IDs** field.
7. From the **Configuration Profile** drop-down list, select a port profile which you want to assign to the specified ports.

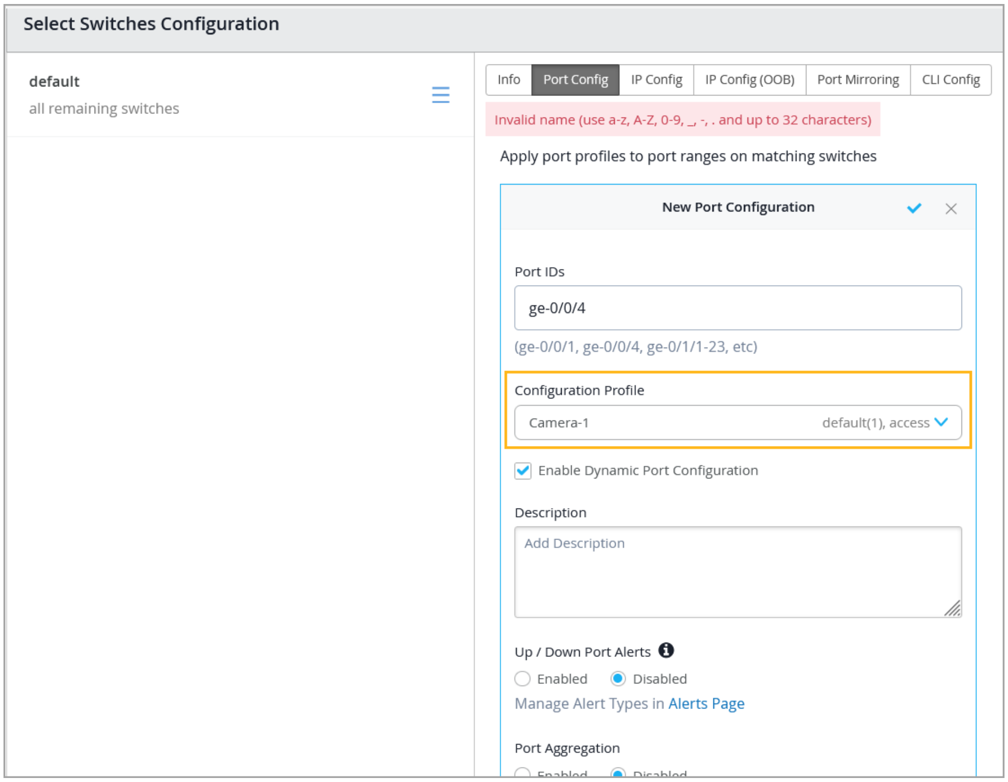

Select Switches Configuration

default

all remaining switches

Invalid name (use a-z, A-Z, 0-9, _, -, . and up to 32 characters)

Apply port profiles to port ranges on matching switches

New Port Configuration ✓ ✕

Port IDs

ge-0/0/4

(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Configuration Profile

Camera-1 default(1), access

☑ Enable Dynamic Port Configuration

Description

Add Description

Up / Down Port Alerts

○ Enabled ● Disabled

Manage Alert Types in Alerts Page

Port Aggregation

○ Enabled ● Disabled

8. Specify other settings as required and then click the check mark (✓) on the upper right of the **New Port Configuration** window.
9. To apply the changes made to the rule, click the check mark (✓) on the upper right of the rule window.
10. Click **Save** on the upper right of the template to confirm and save the changes.

You can also assign port profiles to switch ports at the switch level. You can do that from the **Port** tile in the **Device** section on the switch details page. Click **Add Port Configuration** on the Port tile and specify the settings. Watch the following video for more information:

Video: Port Profiles

Configure Dynamic Port Profile Assignment

Use the steps in this section to configure rules to assign port profiles to an interface dynamically. When a user connects a client device to a switch port with dynamic profile configuration, the switch identifies the device and assigns a suitable port profile to the port. Dynamic port profiling utilizes a set of device properties of the client device to automatically associate a preconfigured port and network setting to the interface. You can configure a dynamic port profile based on the various parameters such as LLDP name and MAC address.

You can configure dynamic port profile assignment at the template level or at the switch level.

To set up dynamic assignment of port profiles at the template level:

1. Go to **Organization** > **Wired** > **Switch Templates** and then click a template to open it.
2. Navigate to the Dynamic Port Configuration tile in the Shared Elements section.
3. On the Dynamic Port Configuration tile, click **Add Rule**.

   The **New Rule** window appears.
4. Set up dynamic port configuration (DPC) rules for automatically assigning port profiles. Here's an example of a rule that automatically assigns the port profile 'AP' to a Mist AP. As per this rule, when the port identifies a device with a chassis ID that starts with D4:20:B0 or D4:21:B1, it assigns the 'AP' profile to the connected device.

   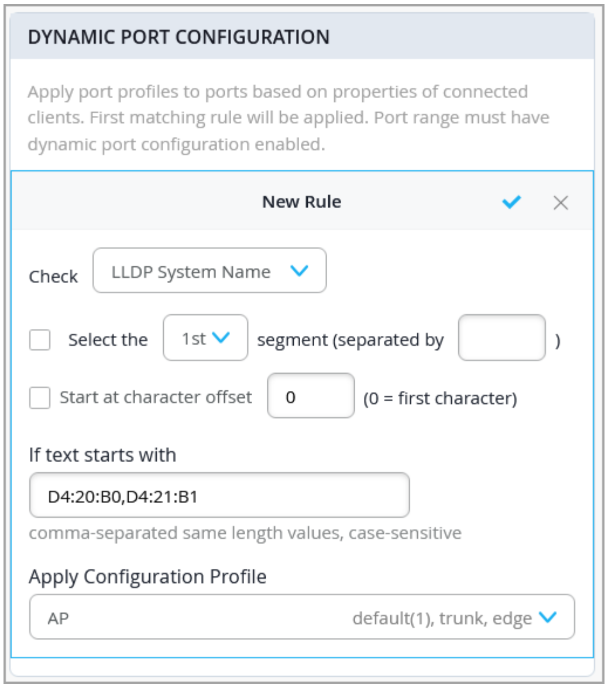

   DYNAMIC PORT CONFIGURATION

   Apply port profiles to ports based on properties of connected clients. First matching rule will be applied. Port range must have dynamic port configuration enabled.

   New Rule ✓ ×

   Check: LLDP System Name

   □ Select the 1st segment (separated by)

   □ Start at character offset 0 (0 = first character)

   If text starts with: D4:20:B0,D4:21:B1

   comma-separated same length values, case-sensitive

   Apply Configuration Profile: AP default(1), trunk, edge

   For more information, refer to the **Dynamic Port Configuration** row in Table 4 on page 49 on the Switch Configuration Options page.

   > **NOTE**: If you use multiple values in the **If text starts with** field in a DPC rule, separate them with commas and ensure that they all have the same length. If any value differs in length, you must create a separate rule for it.

5. To apply the changes, click the check mark (✓) on the upper right of the **New Rule** window.
6. Navigate to the Select Switches Configuration section.
7. Click an existing rule to open it.

   If you want to create a new rule with dynamic port assignment settings, click **Add Rule** on the upper right of the section.
8. Go to the **Port Config** tab of the rule.
9. Click **Add Port Configuration**.

The **New Port Configuration** window appears.

10. Specify the ports in the **Port IDs** field.
11. From the **Configuration Profile** drop-down list, select a port profile. This profile is applied to ports when the connected device does not meet the dynamic assignment rules.

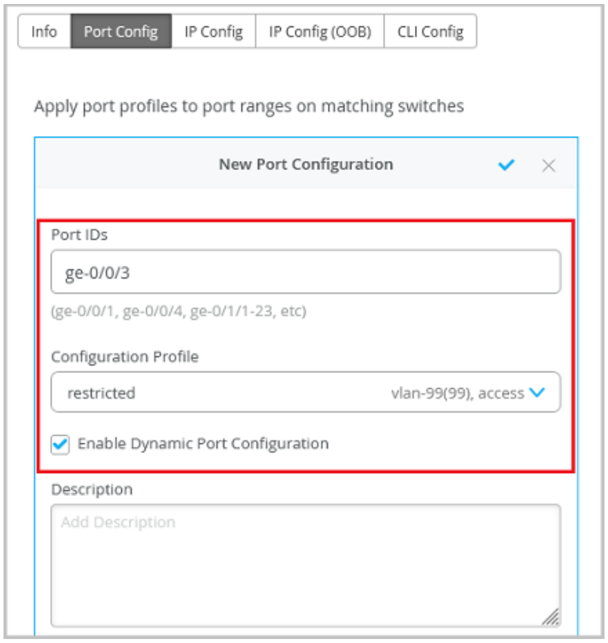

Apply port profiles to port ranges on matching switches

Port IDs: ge-0/0/3

(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Configuration Profile: restricted vlan-99(99), access

☑ Enable Dynamic Port Configuration

Description: Add Description

We recommend that you create a restricted network profile that can be assigned to unknown devices when connected to the switch ports enabled with dynamic port configuration. In the above example, the port is enabled with dynamic port configuration and is assigned with a restricted VLAN. In this case, if the connected device does not match the dynamic profiling attributes, it will be placed into a restricted VLAN such as a non-routable VLAN or a guest VLAN.

12. Select the **Enable Dynamic Configuration** check box.

    For more information, refer to the **Enable Dynamic Port Configuration** row in Table 7 on page 62 on the Switch Configuration Options page.
13. Specify other settings as required and then click the check mark (✓) on the upper right of the **New Port Configuration** window.
14. To apply the changes made to the rule, click the check mark (✓) on the upper right of the rule window.
15. Click **Save** on the upper right of the template to confirm and save the changes.
16. To verify whether a port profile is dynamically assigned to a switch port, follow these steps:

    a. Select **Switches** from the left menu to go to the switch list.

    b. Click the **Switch Insights** link in the Insights column on the switch list.

    c. Look for port assignment events under Switch Events.

You can also set up dynamic port profile assignment at the switch level. You can do that from the switch details page. Watch the following video for more information:

Video: Dynamic Port Profiles (for Colorless Ports)

NOTE:

- Ensure that the default or restricted VLAN used in dynamic port configuration does not have an active DHCP server running. Otherwise, you might encounter stale IP address issue on certain legacy devices.
- A switch with port-based network access control (NAC) authentication does not require dynamic port configuration as VLAN assignments are handled by the RADIUS server. Also, we do not recommend using dynamic port profiles when RADIUS server with MAC Authentication Bypass (MAB) is used.
- Prefer LLDP-based matching over MAC-based matching when the device supports LLDP.
- Do not use MAC-based matching on ports enabled with 802.1X authentication.
- Avoid using `Filter-ID` attributes. When 802.1X is enabled on the ports, VLAN assignment should be handled via RADIUS without relying on `Filter-ID`.

Dynamic port configuration on a switch is meant for establishing connection to IoT devices, APs, and user port endpoints. You should not use it to create connection between switches, switches and routers, and switches and firewalls. You should not enable Dynamic Port Configuration on the uplink port.

When a port profile is assigned to a switch port dynamically based on the connected device, this event is displayed in the Switch Events section on the Switch Insights page. You can also see the dynamic port profile details on a switch port by hovering over the port in the Front Panel section on the switch details page.

NOTE: Junos requires that each interface-range in a port profile contains at least one member interface. When dynamic port configuration is enabled, Junos includes a dummy interface (ge-168/5/X) as a placeholder in the port profile configuration so that the configuration remains valid even when it is not assigned to an actual interface. For instance, if an interface is currently assigned to Port Profile A, but Port Profile B is expected to be dynamically applied later, a placeholder like ge-168/5/0 is used to keep Profile B's interface-range valid.

You can also configure and verify the dynamic port configuration details using the below API:

```
GET - https://api.mistsys.com/api/v1/sites/c3b8f61c-c5a2-495d-9a5c-9b129624e9bf/devices/
00000000-0000-0000-1000-<device mac addr>
```

```
"dynamic": {
            "mode": "dynamic",
            "rules": [
                {
                    "src": "lldp_chassis_id",
                    "usage": "wireless",
                    "equals": "5c:5b:35",
                    "expression": "[0:8]"
                }
            ]
        }

"port_config": {
       "ge-0/0/2-3": {
           "usage": "restricted_access",
           "dynamic_usage": "dynamic"
       }
```

Modify Port Profiles

You can modify port profiles at both the switch template level and the individual switch level.

NOTE:

- You can modify the system defined port profiles only from switch templates. However, Mist does not allow modification of the following system-defined port profiles: **default** and **disabled**.
- If you modify a system-defined port profile from a switch template, that modified version of the profile is available only to that template.

To modify a port profile:

1. Go to **Organization** > **Switch Templates** and the select the appropriate Switch Template.

   If you wish to modify a port profile defined at the switch level, navigate to the switch details page (**Switches** > *switch-name*).

2. Navigate to the **Port Profile** tile.

   In a switch template, you can find this tile in the Shared Elements section. On a switch details page, you can find this tile in the Networks & Port Profiles section.

3. Select the port profile you wish to modify.
4. Make the required changes to the settings as described in Table 5 on page 53.
5. After specifying the settings, click the check mark (✓) on the upper right of the **New Port Profile** window.
6. Click **Save** on the upper right of the template or switch details page to confirm and save the changes.

Delete Port Profiles

You can delete port profiles at the template level or the individual switch level.

To delete a port profile:

1. Go to **Organization** > **Switch Templates** and the select the appropriate Switch Template.

   If you wish to delete a port profile defined at the switch level, navigate to the switch details page (**Switches** > *switch-name*).

   > NOTE:
   >
   > - You can delete the system defined port profiles only from switch templates. However, you cannot delete the following system-defined port profiles: **default** and **disabled**.
   >
   > - From a switch details page, you can only delete the profiles defined at the switch level.

2. Navigate to the **Port Profile** tile.

   In the switch template, you can find this tile in the Shared Elements section. On a switch details page, you can find this tile in the Networks & Port Profiles section.

3. Select the port profile you wish to delete.
4. Click the **trashcan** icon in the top left corner of the **Edit Port Profile** configuration.

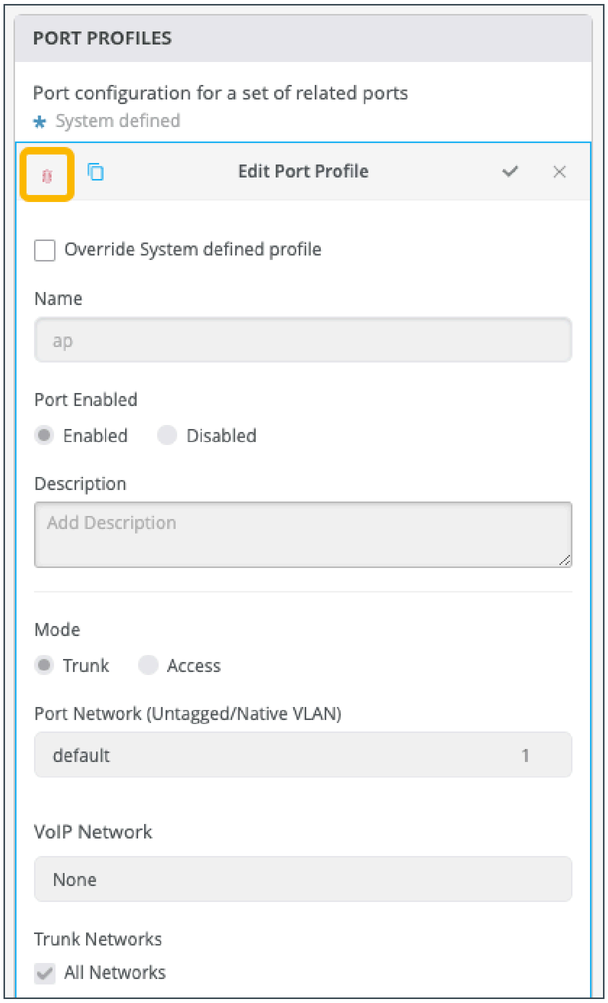

PORT PROFILES

Port configuration for a set of related ports
* System defined

□ Override System defined profile

Name
ap

Port Enabled
● Enabled ○ Disabled

Description
Add Description

Mode
● Trunk ○ Access

Port Network (Untagged/Native VLAN)
default 1

VoIP Network
None

Trunk Networks
☑ All Networks

5. A warning will appear letting you know that the delete action is permanent. You will not be able to recover the port profile once deleted. Enter the name of the port profile, then click **Delete**.

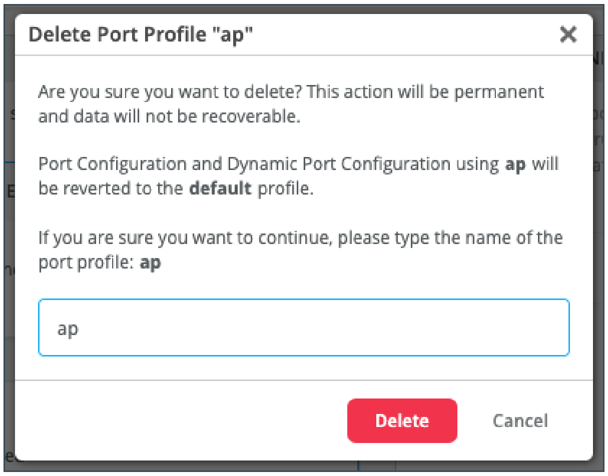

Delete Port Profile "ap"

Are you sure you want to delete? This action will be permanent and data will not be recoverable.

Port Configuration and Dynamic Port Configuration using **ap** will be reverted to the **default** profile.

If you are sure you want to continue, please type the name of the port profile: **ap**

ap

NOTE:

- If you delete the ap, iot, or uplink system-defined port profiles, any reference to these profiles at the Site or device level will revert to the default profile (port configurations or Dynamic Port profiles).
- If you were to create your own port profile and name it “ap”, “iot”, or “uplink” (after having deleted the system-defined port profiles) it will be treated as any other user-defined port profile.

6. Click **Save** on the upper right of the template to confirm and save the changes.

Best Practices in Port Configuration

Here are a few recommendations for your switch ports to work seamlessly with the Mist APs:

- On a trunk port, prune all the unwanted VLANs. Only the required VLANs (based on the WLAN configuration) should be on the port. Since the APs do not save the configuration by default, APs should be able to get the IP address on the native VLAN to get connected to the cloud and get configured.
- We do not recommend port security (MAC address limit), except in the case where all WLANs are tunneled.
- Feel free to enable BPDU guard, as BPDUs are typically not bridged from wireless to wired connection on an AP unless it is a mesh base. **BPDUs** are data messages that are exchanged across the switches within an extended LAN that uses a spanning tree protocol topology. **BPDU** packets contain information on ports, addresses, priorities, and costs and ensure that the data ends up where it was intended to go.

Enable Port Mirroring

Port mirroring copies source traffic traversing one or more interfaces and forwards it at line rate to a specified destination. It is useful for monitoring network traffic, enforcing network policies, identifying network issues, and as a part of intrusion detection.

You can set up port mirroring for ingress and/or egress traffic on any interface on the switch (except those identified below), as well as for ingress traffic on a VLAN (*network*, in the terminology of Mist dashboard), private VLAN, or bridge domain. You can apply it to all traffic on the port, or a filtered subset of the traffic, as defined in a firewall policy. Mirrored traffic is sent to a local interface, VLAN, or bridge domain for remote monitoring or analysis.

From the Mist dashboard, you can configure port mirroring for managed switches that are part of a switch template, as well as for independent switches. To do so, you'll need to have an *input*, that is, the source of the traffic, and an *output*, that is, the destination mirrored traffic will be forwarded to. For switch ports, you can specify whether you want to mirror traffic in the *ingress* direction or *egress*. If you select a VLAN as the source, *ingress* mirroring is assumed.

Mirroring is not supported for management interfaces (such as me0 or vme0), fiber channel interfaces, or integrated routing and bridging (IRB) interfaces (which are also known as routed VLAN interfaces or RVIs). A maximum of four port mirroring configurations is allowed.

NOTE: To specify a VLAN in the port mirroring configuration, select **network** and choose from the list of VLANs available in Mist. To specify a bridge domain, select the **IP address** option.

To set up port mirroring on a managed switch that is part of a switch template:

1. Select **Switches** from the main menu and then select the switch you want to configure from the list of switches that appears.
2. In the switch configuration page that opens, scroll down to find the **Port Mirroring** section.
3. Since we are configuring a switch that belongs to a template in this example, select **Override Site/ Template Settings** to enable the **Add Port Mirror** option. (If you are setting up port mirroring in a switch template or on an unmanaged switch, **Add Port Mirror** is already enabled.)
4. Click **Add Port Mirror** to open the **Edit Port Mirror** screen, as shown in the figure below.

Figure 3: Configure Port Mirroring in Mist

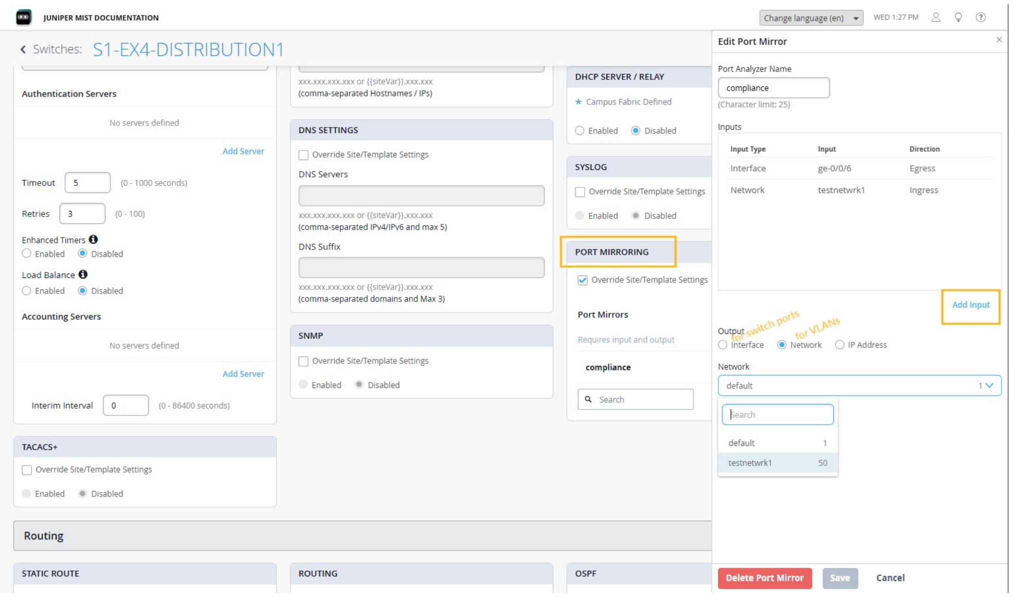

Timeout 5 (0 - 1000 seconds)

Retries 3 (0 - 100)

Enhanced Timers ○ Enabled ● Disabled

Load Balance ○ Enabled ● Disabled

Interim Interval 0 (0 - 86400 seconds)

TACACS+: □ Override Site/Template Settings ○ Enabled ● Disabled

xxx.xxx.xxx.xxx or {{siteVar}}.xxx.xxx (comma-separated Hostnames / IPs)

DNS SETTINGS: □ Override Site/Template Settings

xxx.xxx.xxx.xxx or {{siteVar}}.xxx.xxx (comma-separated IPv4/IPv6 and max 5)

xxx.xxx.xxx.xxx or {{siteVar}}.xxx.xxx (comma-separated domains and Max 3)

SNMP: □ Override Site/Template Settings ○ Enabled ● Disabled

DHCP SERVER / RELAY: ○ Enabled ● Disabled

SYSLOG: □ Override Site/Template Settings ○ Enabled ● Disabled

PORT MIRRORING: ☑ Override Site/Template Settings

Edit Port Mirror

Port Analyzer Name: compliance (Character limit: 25)

Inputs

Input Type	Input	Direction
Interface	ge-0/0/6	Egress
Network	testnetwrk1	Ingress

Output ○ Interface ● Network ○ IP Address

Network: default 1

Search	
default	1
testnetwrk1	50

5. Give the port mirroring instance a name in the **Port Analyzer Name** field.
6. Click **Add Input** and then choose from the following:
    - **Interface--**to specify a local port on the switch as the source of the traffic to mirror. You can also click the icon to see the which ports are configured on the switch and select one, or multiple, ports from the interactive map that pops up.
        - **Ingress-**to mirror traffic that is inbound to the switch interface.

          Egress--to mirror traffic that is outbound from the switch interface.
    - **Network-**to specify a VLAN as the traffic source (inbound traffic only).
7. Click the blue check box to accept your input configuration, before moving on to add a destination for the mirrored traffic under **Output**. Choose from the following:
    - **Interface--**to specify a local port on the switch as the destination of the mirrored traffic. You can also click the icon to see the which ports are configured on the switch and select one of those.
    - **Network-**to specify an existing VLAN as the destination for the mirrored traffic.
    - **IP address-**to specify the IP address as the destination for the mirrored traffic.
8. Click the blue check box to accept your configuration, then **Add** post your configuration changes, and **Save**, in the upper right corner of the page, to actually commit your changes on the switch.

To confirm that port mirroring is enabled on the switch, click **Switches** in the main menu and then **Switch Insights** for the switch you just configured. The update should appear within a few minutes under **Switch Events** section of the dashboard.

Configure Port Speed

SUMMARY

Use the steps in this topic to configure port channelization or port speed.

You can channelize ports and configure chassis speed (port speed) on a select set of switches. Port channelization enables you to split a high-speed QSFP or QSFP+ port on a switch into several low-speed ports. For example, you can channelize a 100 Gbps port into four 25 Gbps ports.

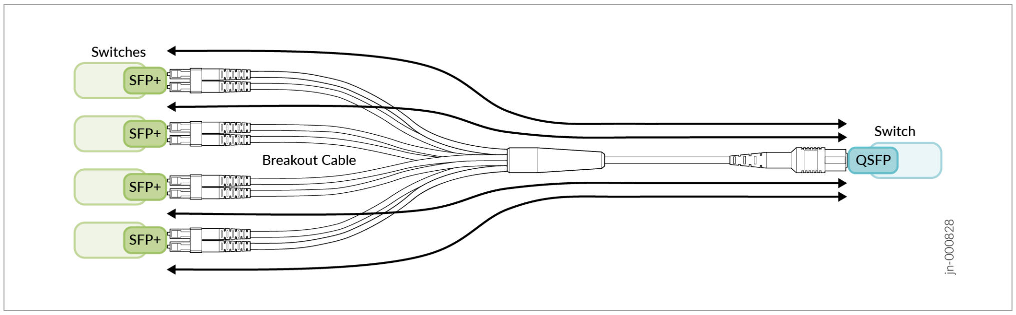

Chassis speed configuration is available for SFP or SFP+ ports. Note that the SFP and SFP28 ports are grouped in quads (groups of four), and you can configure the speed of the ports only in quads; you cannot configure the speed for a single SFP28 port. The speed set for the first port in a quad gets applied to all the ports in that quad.

The following table lists the devices and their respective ports that support port channelization, along with the supported channel speeds:

Model	Port	Channel Speed
QFX5120-32C	0-30	4x10g, 4x25g, 2x50g
	31	2x50g
QFX5120-48T	48-49, 52-53	2x50g
	50, 51	4x10g, 4x25g, 2x50g
QFX5120-48Y	48-55	4x10g, 4x25g

(Continued)

Model	Port	Channel Speed
QFX5120-48YM	50, 52	4x10g, 4x25g, 2x50g
QFX5130-32CD	0-31	4x10g, 4x25g, 2x50g, 2x100g, 8x50g, 4x100g, 2x200g
EX4650-48Y	48-55	4x10g, 4x25g

The following table lists the devices and their respective ports that support chassis speed configuration, along with the supported speeds:

Model	Port	Chassis Speed
QFX5120-48T	0-47	1g, 10g (per quad 0, 4, 8 .. . 44)
QFX5120-48Y	0-47	1g, 10g, 25g (per quad 0, 4, 8 .. . 44)
QFX5120-48YM	0-47	1g, 10g, 25g (per quad 0, 4, 8 .. . 44)
EX4650-48Y	0-47	1g, 10g, 25g (per Quad 0, 4, 8 .. . 44)

You can view the channelization information for a port by hovering on it on the front panel on the switch details page.

To channelize ports or configure chassis speed on a switch:

1. In the Mist portal, click **Switches** > *Switch Name* to navigate to the switch details page.
2. Locate the Port tile and click the **Advanced** tab.

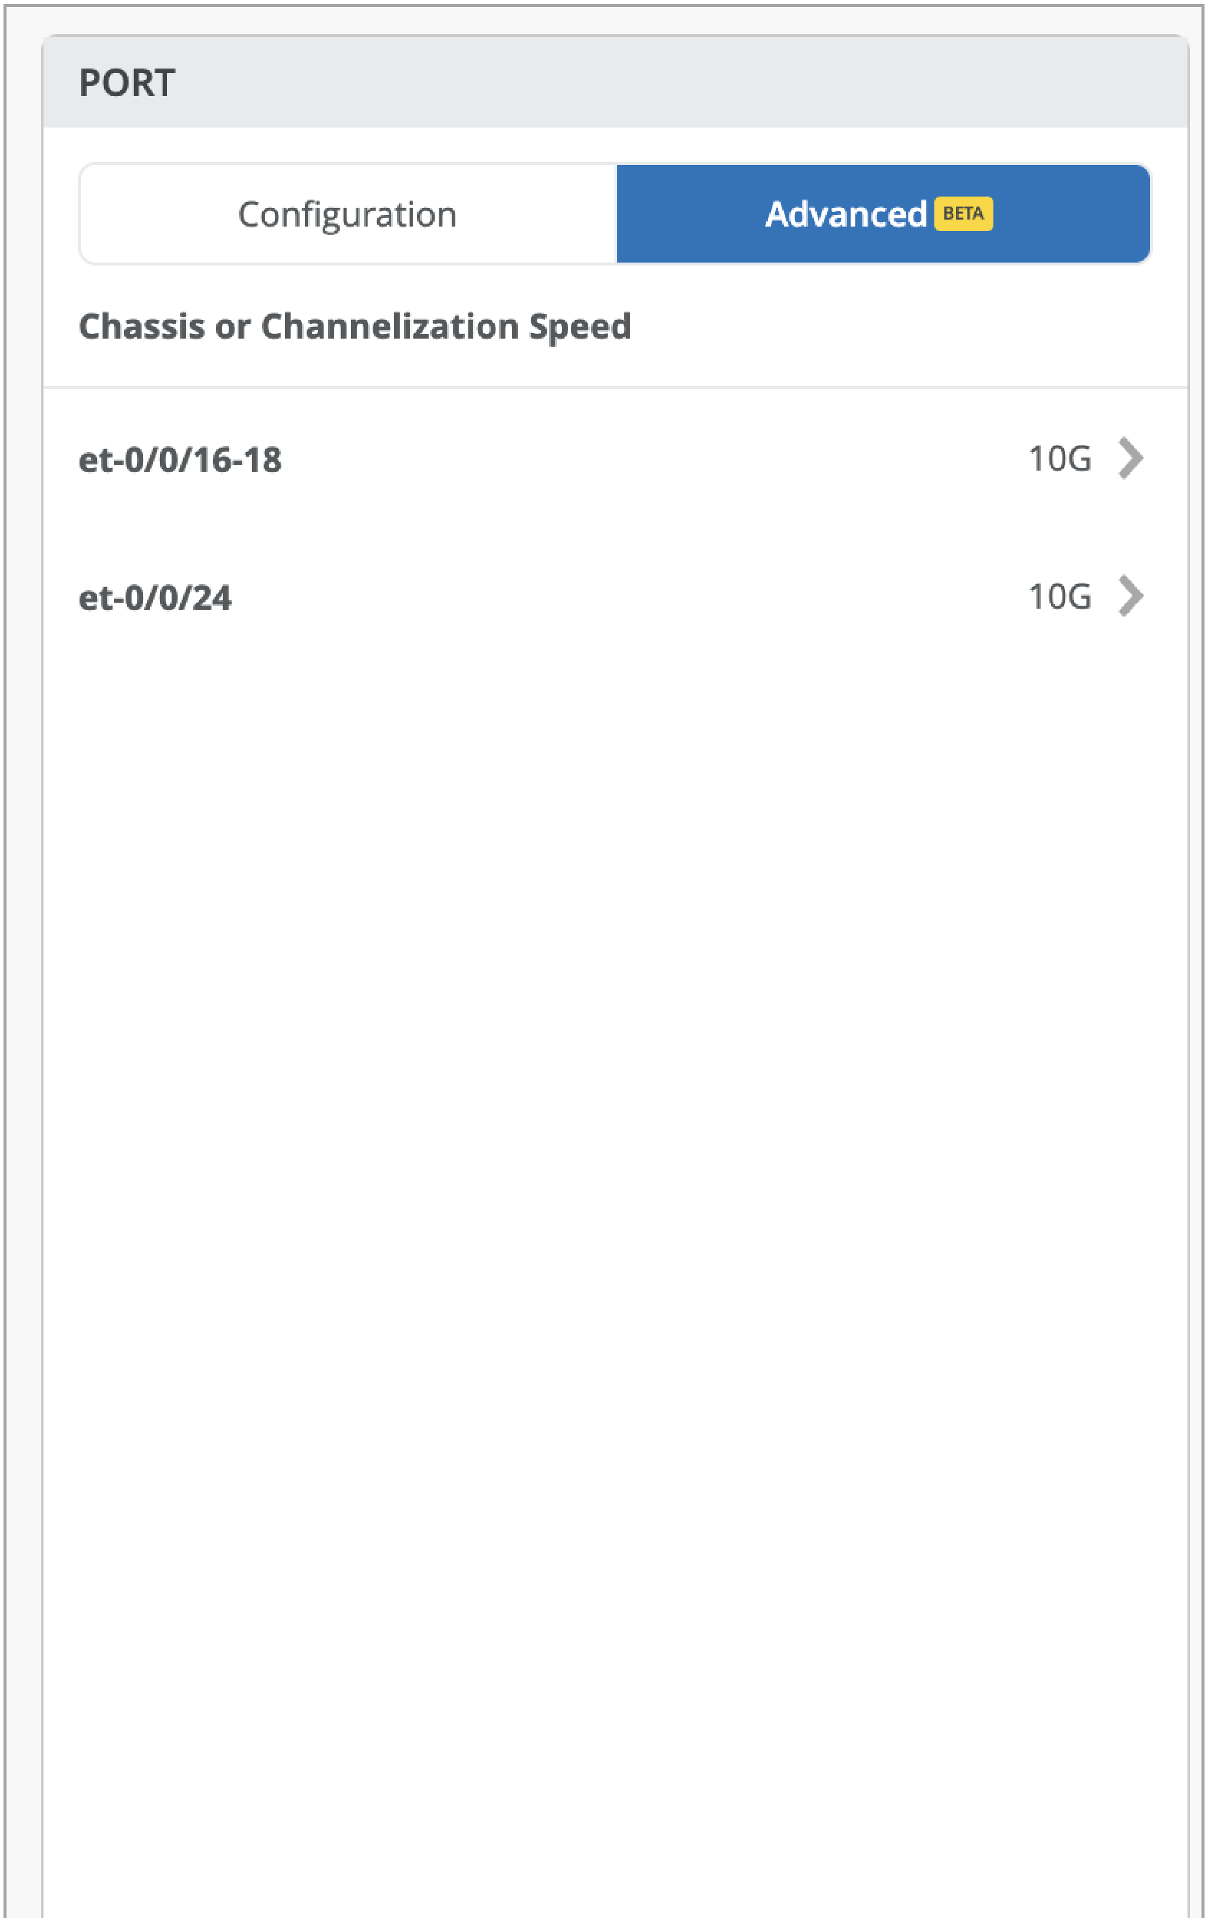

PORT

Chassis or Channelization Speed

et-0/0/16-18	10G
et-0/0/24	10G

3. Click **Advanced Port Configuration**.
4. Specify the following:
   - **Port IDs**—Use the port selector to specify the ports on which you need to configure the chassis speed or channel speed.
   - **Chassis or Channel Speed**—Specify the chassis speed or channel speed here. Chassis speed applies only to SFP and SFP+ ports. If you selected QSFP or QSFP+ ports, configuring a speed channelizes the port, enabling you to split the port into several low-speed ports. Available speeds vary depending on the switch model selected.
5. Click **Add**.

Once you channelize a port, you can assign port profiles to individual channel speeds.

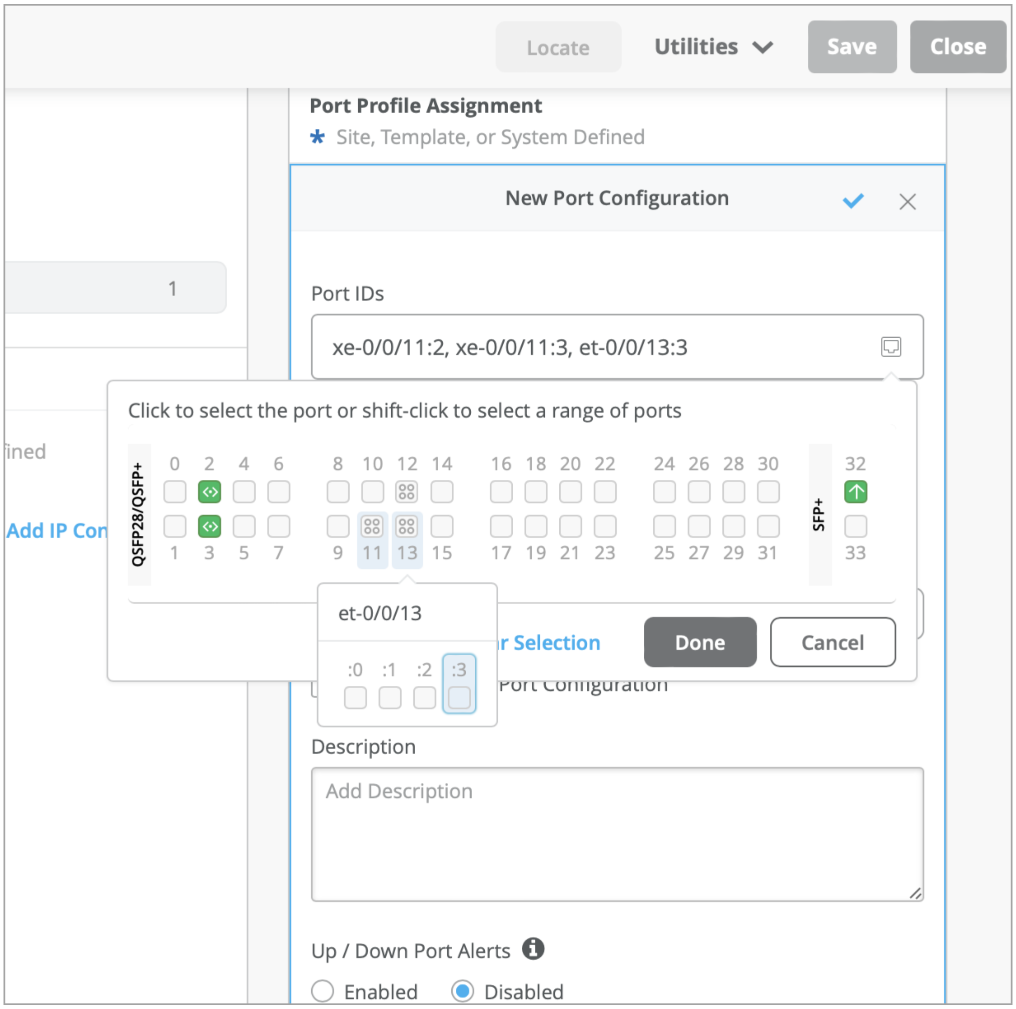

Port Profile Assignment

* Site, Template, or System Defined

New Port Configuration

Port IDs: xe-0/0/11:2, xe-0/0/11:3, et-0/0/13:3

Click to select the port or shift-click to select a range of ports

Description: Add Description

Up / Down Port Alerts

○ Enabled ● Disabled

Similarly, the campus fabric configuration screen allows you to select the channelized ports while configuring EVPN connections.

To delete a port speed configuration, navigate to the **Advanced** tab of the Port tile on the switch details page and click the configuration to be deleted and then click **Delete Port Config**.

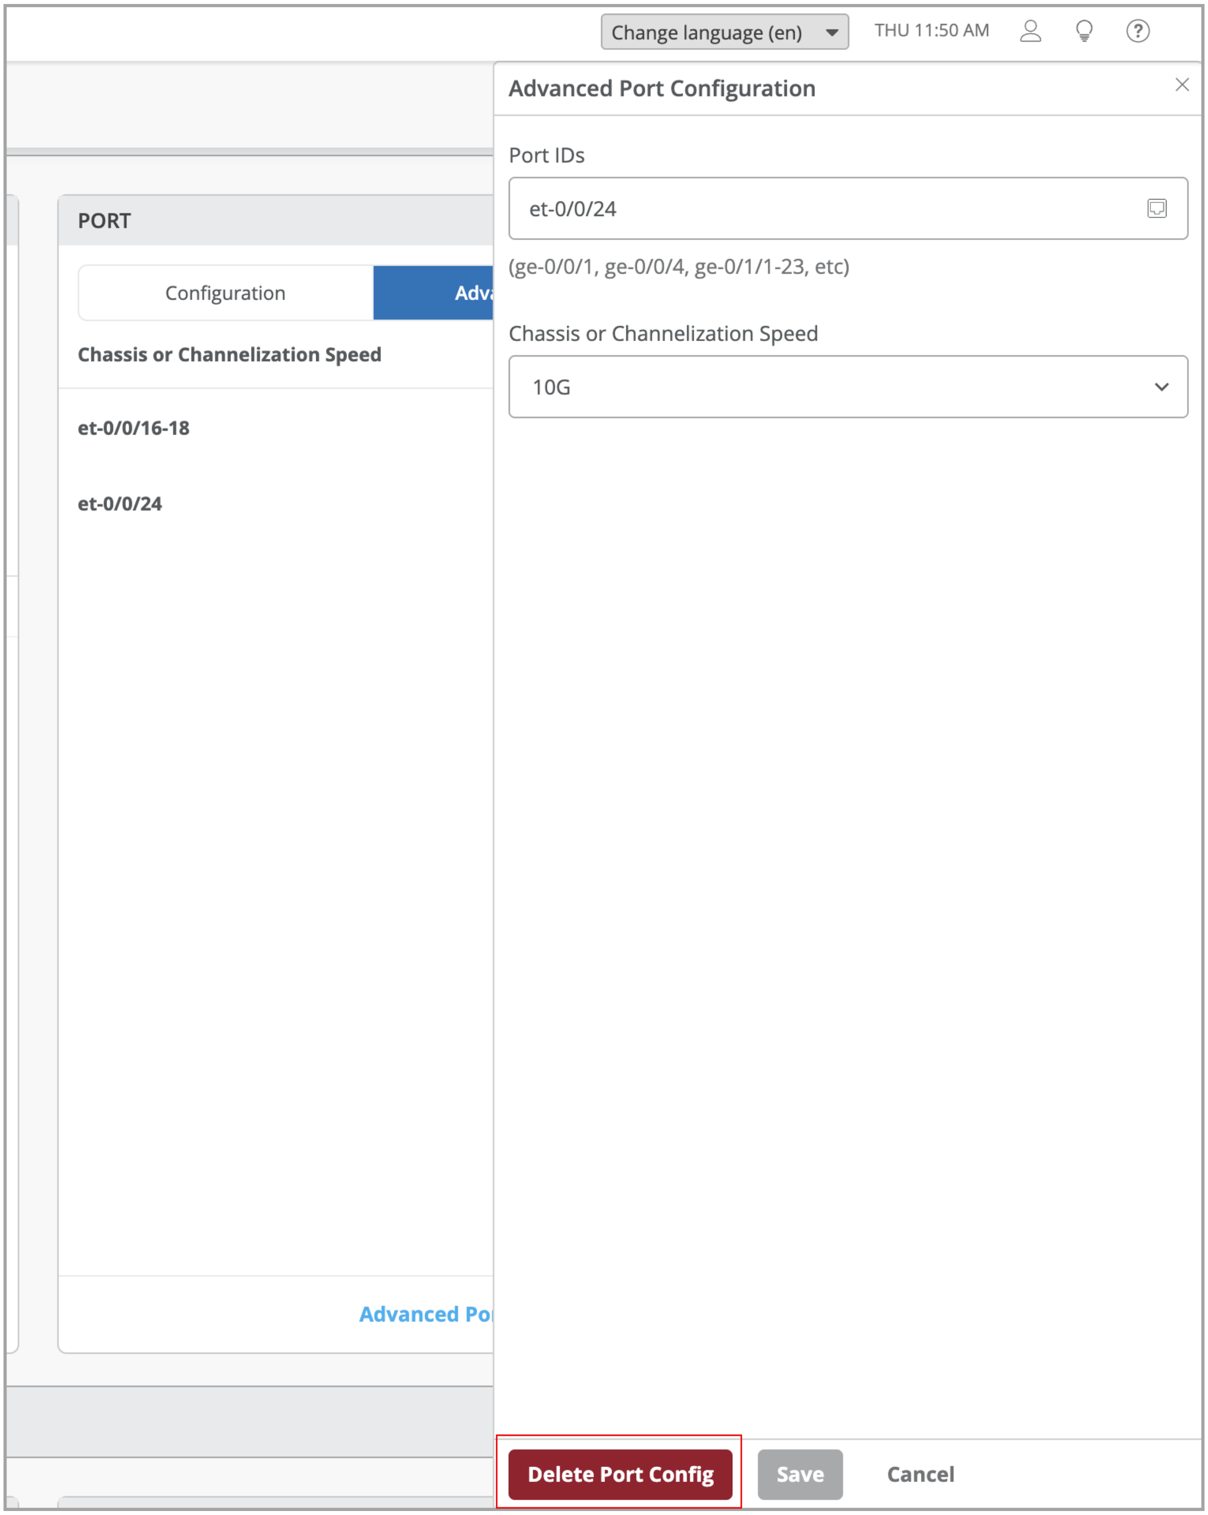

Advanced Port Configuration

Port IDs: et-0/0/24

(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Chassis or Channelization Speed: 10G

Configure BGP on Switches via Mist

SUMMARY

Understand the benefits of Border Gateway Protocol (BGP) and follow these steps to configure a BGP group on your switch.

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routers or layer 3 switches in different autonomous systems (AS). The routing information includes the complete route to each destination. BGP uses this information to maintain a database of network reachability information, which it exchanges with other BGP systems. BGP uses the network reachability information to construct a graph of AS connectivity, which enables BGP to prevent routing loops and enforce policy decisions at the AS level.

At a high level BGP configuration includes a BGP group, BGP neighbors, and a routing policy to advertise the BGP routes.

For BGP configuration examples and CLI steps, see BGP Configuration Overview.

Before you configure BGP, ensure that you have configured the network interfaces (L3 interfaces). You can configure L3 interface from the Port Configuration tile on the switch details page (**Switches** > **Switch Name**). In this configuration, you need to specify the ports, interface type (L3), IP address, and subnet mask (see the image below).

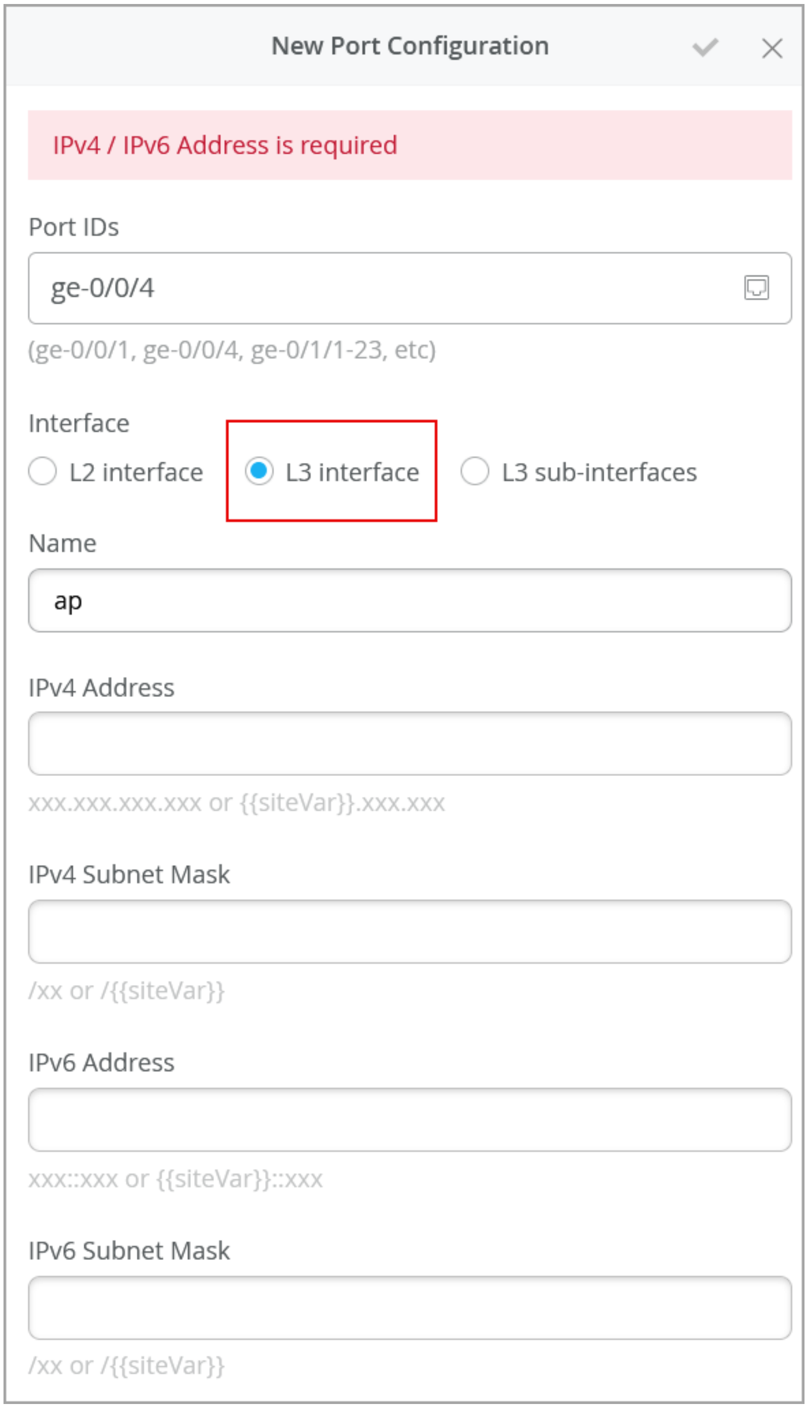

New Port Configuration

IPv4 / IPv6 Address is required

Port IDs

ge-0/0/4

(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Interface

○ L2 interface ● L3 interface ○ L3 sub-interfaces

Name

ap

IPv4 Address

xxx.xxx.xxx.xxx or {{siteVar}}.xxx.xxx

IPv4 Subnet Mask

/xx or /{{siteVar}}

IPv6 Address

xxx::xxx or {{siteVar}}::xxx

IPv6 Subnet Mask

/xx or /{{siteVar}}

To configure BGP on a switch via Mist:

1. In the Mist portal, click **Switches** to go to the list of switches.
2. From the switch list, click the switch on which you want to configure BGP. The switch details page is displayed.
3. In the BGP section on the switch details page, select the **Enabled** check box.
4. Click **Add BGP Group**.

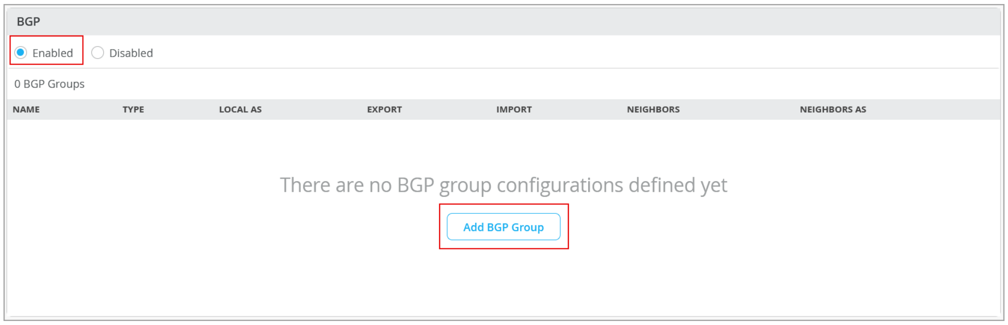

BGP

● Enabled ○ Disabled

0 BGP Groups

NAME	TYPE	LOCAL AS	EXPORT	IMPORT	NEIGHBORS	NEIGHBORS AS

There are no BGP group configurations defined yet

The Add BGP Group window is displayed.

5. Configure the BGP Group settings as described below:

Field	Description
Name	Specify a name for the BGP group.
Type	Select one of the following BGP type: • Internal: When two BGP-enabled devices are system (AS), the BGP session is called an inte session. IBGP routes traffic within an autono • External: External BGP routes traffic betweer different autonomous systems.
Network (VLAN)	Select a network. The BGP group will be added t is a part of. If the VRF has more than one networ selected.
BFD Interval (in milliseconds)	Specify the Bidirectional Forwarding Detection (I milliseconds. BFD protocol is a simple hello mech in a network. Hello packets are sent at a specifie neighbor failure is detected when the routing de after a specified interval.
Local AS	Specify the local autonomous system (AS) numbe see Understanding the BGP Local AS Attribute.

Hold Time	Specify the hold-time value to use when negotia peer. The hold-time value is advertised in open p peer the length of time that it should consider th does not receive a keepalive, update, or notificat specified hold time, the BGP connection to the p devices through that peer become unavailable. The device calculates the BGP keepalive interval configured hold time. For instance, if you set the the keepalive interval will be set to 30 seconds.
Authentication Key	Configure an MD5 authentication key (password devices use the same password to verify the autl sent from this system.
Export	Select a routing policy to determine which route peers. An export policy is applied to outbound ro output of the `show route advertising-protocol bg` command. If you don’t have a policy available, click Create F down menu to create a new one. For more inforr Routing Policies on Switches via Mist" on page 9
Import	Select a routing policy to determine which route peers and added to the local routing table. An im inbound routes that are visible in the output of tl `protocol bgp neighbor-address` command. If you don’t have a policy available, click Create F down menu to create a new one. For more inforr Routing Policies on Switches via Mist" on page 9

6. Click **Add Neighbor**.

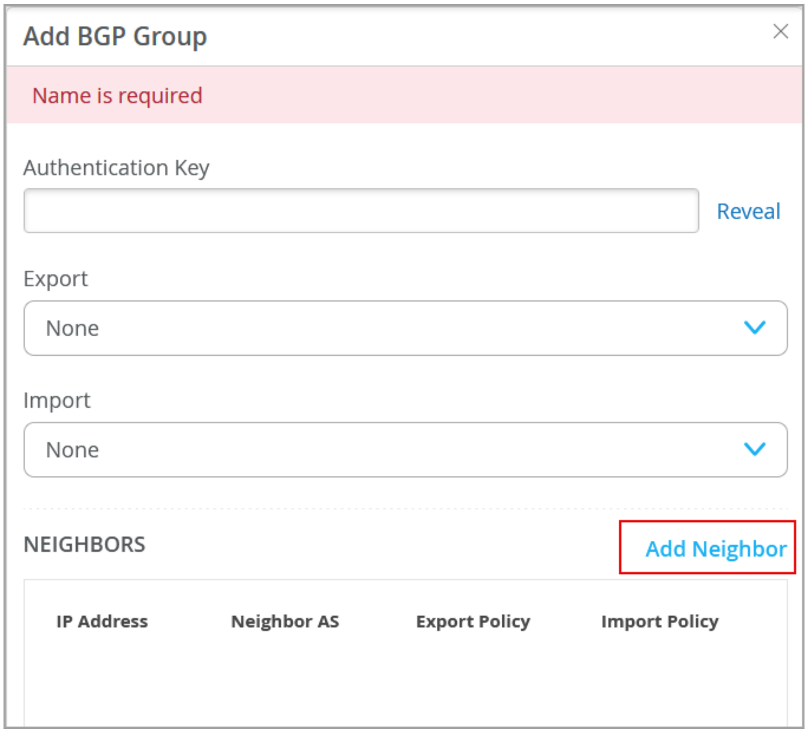

Add BGP Group

Name is required

Authentication Key

Reveal

Export

None

Import

None

NEIGHBORS

Add Neighbor

IP Address	Neighbor AS	Export Policy	Import Policy

The Add Neighbor (IPv4 / IPv6) window is displayed.

7. Configure the neighbor information as described below:

Field	Description
IP Address	Specify the IP address of the remote BGP peer.
Neighbor AS	Specify the AS number of the remote BGP peer.
Hold Time	Specify the hold-time value to use when negotia peer.
Export	Select a routing policy to determine which route peer.
Import	Select a routing policy to determine which route BGP peer and added to the local routing table.

8. After configuring the neighbor information, click the check mark on the upper right of the Add Neighbor window.
9. Click the **Add** button on the lower left of the Add BGP window.
10. To save the configuration, click **Save** on the upper right of the switch details page.

Configure Routing Policies on Switches via Mist

SUMMARY

Understand the benefits of routing policies, and follow these steps to add one to a switch template.

Routing policies allow you to control the routing information between the routing protocols and the routing tables and between the routing tables and the forwarding table. All routing protocols use the Junos OS routing tables to store the routes that they learn and to determine which routes they should advertise in their protocol packets. Routing policies also allow you to control which routes the routing protocols store in and retrieve from the routing table.

The routing policy is composed of terms. Each term can include a set of conditions and a then statement, which defines the actions to take if a route matches the conditions specified in the term.

You can configure routing policy at the organization level (**Organization** > **Switch Templates**), site level (**Site**> **Switch Configuration**) or at the switch level (**Switches** > ***Switch Name***). Ideally you don't need to configure a routing policy per switch.

To configure a routing policy at the organization level:

1. Click **Organization** > **Switch Templates** to go to the list of switch templates.
2. Click the switch template associated to the site where you want to configure the routing policy. The switch template is displayed.
3. In the Routing Policy section on the switch template, click **Add Routing Policy**.

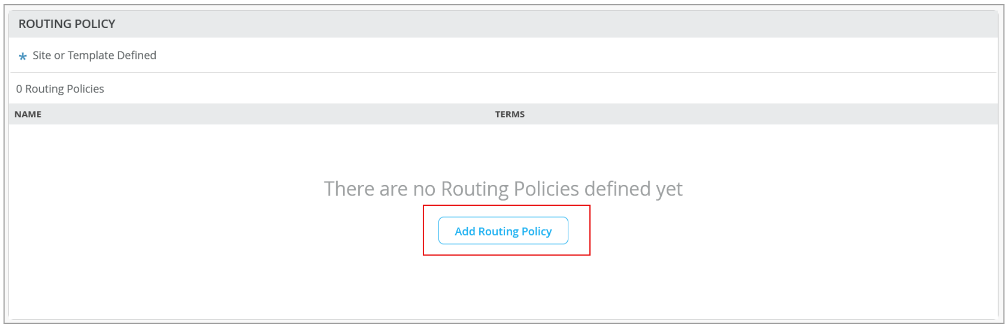
ROUTING POLICY
* Site or Template Defined
0 Routing Policies
NAME
TERMS
There are no Routing Policies defined yet
Add Routing Policy

The Create Routing Policy window is displayed.

NOTE: If you want to configure routing policy the switch level, go to the Routing Policy section on the switch details page (**Switches** > ***Switch Name***). Similarly, if you want to configure routing policy the site level, go to the Routing Policy section on the site template (**Site** > **Switch Configuration**).

4. Provide a name for the routing policy in the **Name** field.
5. Click **Add Terms**.
   The Add Terms section is displayed.
6. Configure the terms as described below:

Field	Description
Name	Specify a name for the term.
Prefix	Specify a prefix against which the incoming routes destination prefix with mask that is greater than or prefix length, use a range of masks. This field supp Example: 10.10.10.0/24-32 would allow all prefixe
AS Path	Specify the AS path. A BGP AS path is the sequenc that network packets traverse to get to a specified information, see Understanding AS Path Regular E Routing Policy Match Conditions.
Protocol	Select a routing protocol.
Community	Specify a BGP community. A BGP community is a share a common property. Community information attribute in BGP update messages.
Then	Define an action to be applied on the route, if it m specified. Options are: • Accept • Reject

Add Action	You can further specify one or more additional act route characteristics. The following actions are ava • Set Community: (BGP only) This action replace were in the route in with the specified commur • Prepend AS Path: (BGP only) This action affixes at the beginning of the AS path. • Set Local Preference: (BGP only) This action se preference attribute. The preference value can from 0 through 4,294,967,295.

7. After configuring the terms information, click the check mark on the upper right of the Add Terms section on the Create Routing Policy window.
8. Click the **Add** button on the lower left of the Create Routing Policy window.
9. To save the configuration, click **Save** on the upper right of the switch template.

Protection of Routing Engine

IN THIS SECTION

The Protection of Routing Engine feature ensures that the Routing Engine accepts traffic only from trusted systems. Enabling this feature results in creation of a stateless firewall filter that discards all traffic destined for the Routing Engine, except those from the specified trusted sources. Protecting the Routing Engine involves filtering incoming traffic on the router's lo0 interface. Enabling this feature on Juniper Switches is suggested as a best practice.

Configure Protection of Routing Engine

When Protection of Routing Engine is enabled, Mist by default ensures that the following services (if configured) are allowed to communicate with the switch: BGP, BFD, NTP, DNS, SNMP, TACACS, RADIUS, and Mist cloud connectivity.

If you want to additionally configure ICMP or SSH to access the switch from, you can enable them under Trusted Services. Note that enabling ICMP and SSH opens these protocols to all networks.

If you want to configure the commonly used IP networks to access the switch from, you can configure that under Trusted Networks. Use this option if you want to access the switch from the entire network.

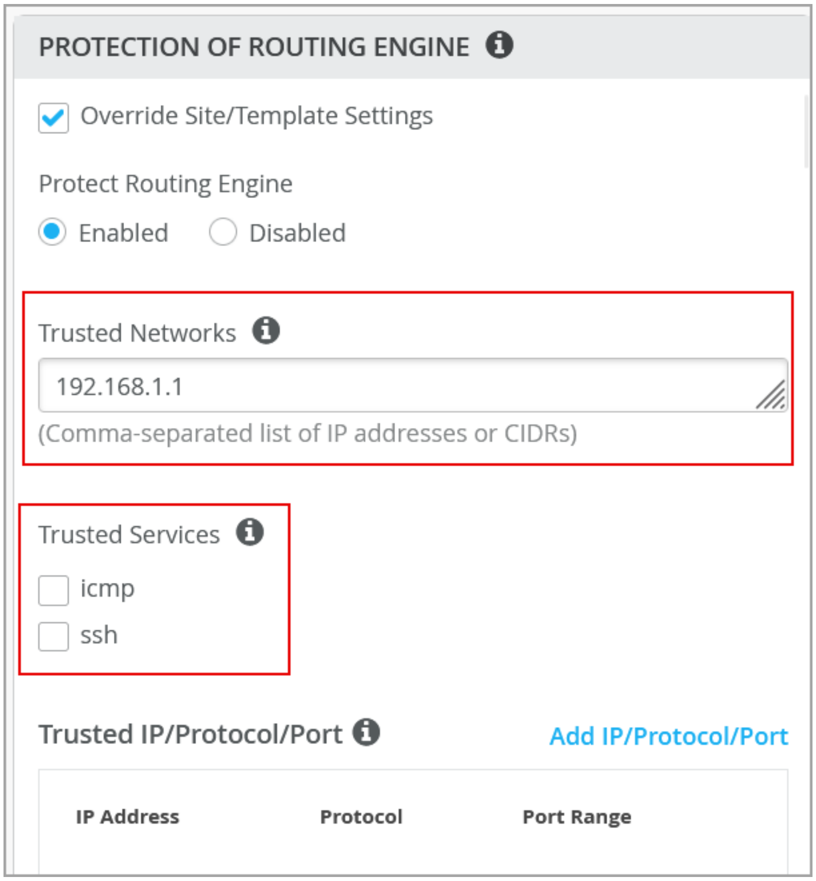

PROTECTION OF ROUTING ENGINE

☑ Override Site/Template Settings

Protect Routing Engine

● Enabled ○ Disabled

Trusted Networks

192.168.1.1

(Comma-separated list of IP addresses or CIDRs)

Trusted Services

□ icmp

□ ssh

Trusted IP/Protocol/Port

IP Address	Protocol	Port Range

If you have other custom services (which are a specific combination of IP, Port and Protocol) that you would like to reach the switch from, you can configure them under Trusted IP/Port/Protocol. This option allows you to use a particular port and protocol to access the switch.

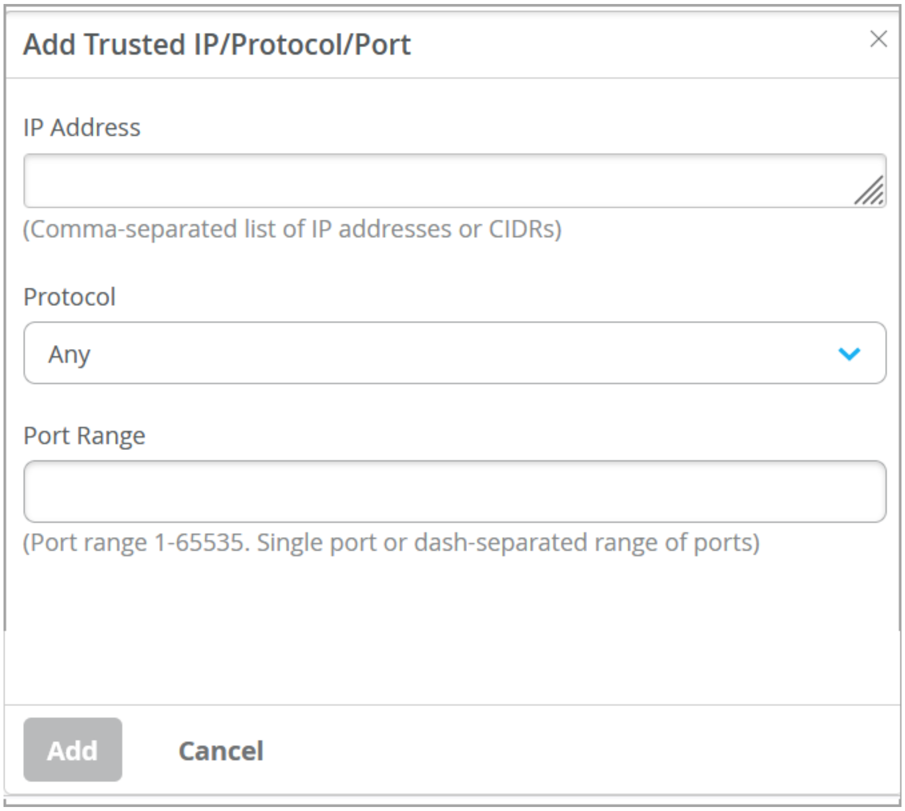

Add Trusted IP/Protocol/Port

IP Address

(Comma-separated list of IP addresses or CIDRs)

Protocol

Any

Port Range

(Port range 1-65535. Single port or dash-separated range of ports)

You can configure Protection of Routing Engine at the organization level (Organization > Switch Templates), at the site level (Site > Switch Configuration), and at the switch level (Switches > Switch Name).

The following procedure lists steps for configuring Protection of Routing Engine at the switch level.

To configure Protection of Routing Engine at the switch level:

1. Click **Switches** > *switch name* to navigate to the switch details page.
2. Scroll down to the **PROTECTION OF ROUTING ENGINE** tile in the Management section.
3. Select the **Override Site/Template Settings** check box.
4. Select the **Enabled** check box.

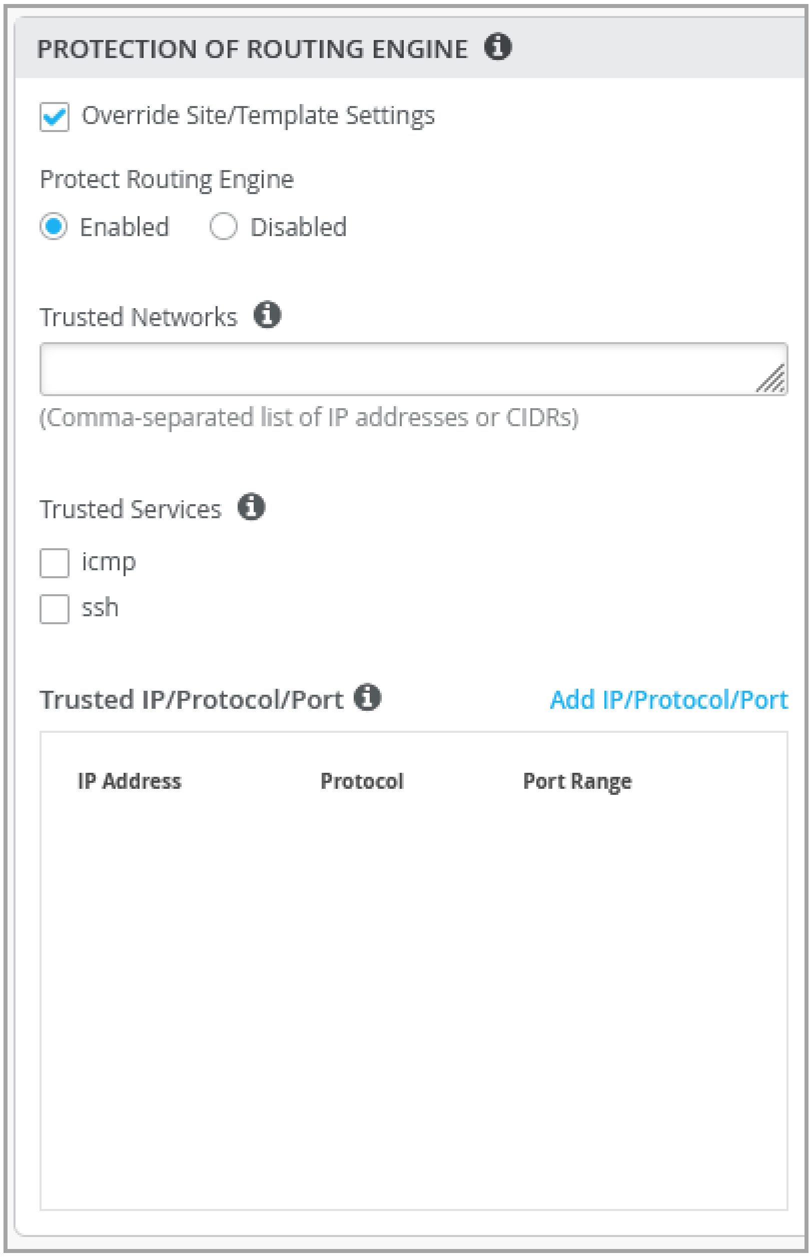

PROTECTION OF ROUTING ENGINE

☑ Override Site/Template Settings

Protect Routing Engine

● Enabled ○ Disabled

Trusted Networks

(Comma-separated list of IP addresses or CIDRs)

Trusted Services

□ icmp

□ ssh

Trusted IP/Protocol/Port

IP Address	Protocol	Port Range

When Protection of Routing Engine is enabled, Mist automatically parses the configuration and allows the end hosts (BGP neighbors, DNS/NTP/TACACS/RADIUS servers, SNMP Clients etc) to communicate with the switch. If you want to add additional IP or IP Subnet that you want the switch to communicate with, add those networks in the Trusted Networks section as mentioned in the next step.

5. To add additional IP or IP Subnet that you want the switch to communicate with, enter the IP addresses in a comma separated format in the **Trusted Networks** field.

6. If you want the switch to respond to the SSH and ICMP services, select the **ssh** and **icmp** check boxes.

7. If you want the switch to respond to custom services (which are a specific combination of IP, Port and Protocol), follow the below steps:

   a. Click **Add IP/Protocol/Port**.
      The Add Trusted IP/Protocol/Port window is displayed.

   b. In the Add Trusted IP/Protocol/Port window, specify the IP Address, a Protocol, and an applicable Port Range.

   c. Click **Add**.

8. Save the configuration.

Configuration Commands (CLIs)

```
"set groups top firewall family inet filter protect_re term allow_mist_obssh from source-port
[ 2200 ]",
"set groups top firewall family inet filter protect_re term allow_mist_obssh then accept",
"set groups top firewall family inet filter protect_re term allow_dhcp from source-port [ 67
68 ]",
"set groups top firewall family inet filter protect_re term allow_dhcp from destination-port
[ 67 68 ]",
"set groups top firewall family inet filter protect_re term allow_dhcp from protocol udp",
"set groups top firewall family inet filter protect_re term allow_dhcp then accept",
"set groups top firewall family inet filter protect_re term allow_bgp from source-prefix-list
bgp_neighbors",
"set groups top firewall family inet filter protect_re term allow_bgp from source-prefix-list
bgp_vrf_neighbors",
"set groups top firewall family inet filter protect_re term allow_bgp from destination-port
[ 179 ]",
"set groups top firewall family inet filter protect_re term allow_bgp from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_bgp then accept",
"set groups top firewall family inet filter protect_re term allow_bfd from source-prefix-list
bgp_neighbors",
"set groups top firewall family inet filter protect_re term allow_bfd from source-prefix-list
bgp_vrf_neighbors",
"set groups top firewall family inet filter protect_re term allow_bfd from destination-port
[ 3784 4784 ]",
"set groups top firewall family inet filter protect_re term allow_bfd from protocol udp",
"set groups top firewall family inet filter protect_re term allow_bfd then accept",
"set firewall family inet filter protect_re term allow_ntp_src from protocol udp",
"set firewall family inet filter protect_re term allow_ntp_src from source-port 123",
"set firewall family inet filter protect_re term allow_ntp_src then accept",
"set firewall family inet filter protect_re term allow_ntp_dest from protocol udp",
"set firewall family inet filter protect_re term allow_ntp_dest from destination-port 123",
"set firewall family inet filter protect_re term allow_ntp_dest then accept",
"set groups top firewall family inet filter protect_re term allow_dns from source-port [ 53 ]",
"set groups top firewall family inet filter protect_re term allow_dns from protocol [ tcp
udp ]", "set groups top firewall family inet filter protect_re term allow_dns then accept",
"set groups top firewall family inet filter protect_re term allow_radius from source-prefix-list
radius_servers",
"set groups top firewall family inet filter protect_re term allow_radius from destination-port
[ 1812 1813 ]",
"set groups top firewall family inet filter protect_re term allow_radius from protocol udp",
"set groups top firewall family inet filter protect_re term allow_radius then accept",
"set groups top firewall family inet filter protect_re term allow_tacacs from source-prefix-list
tacacs_servers",
"set groups top firewall family inet filter protect_re term allow_tacacs from destination-port
[ 49 ]",
"set groups top firewall family inet filter protect_re term allow_tacacs from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_tacacs then accept",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from source-
prefix-list snmp_clients",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from destination-
port [ 161 10161 ]",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from protocol
udp",
"set groups top firewall family inet filter protect_re term allow_snmp_clients then accept",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 10-216-192-1_32",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 100-100-100-2_32",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 8-8-8-8_32",
"set groups top firewall family inet filter protect_re term trusted_hosts then accept",
"set groups top firewall family inet filter protect_re term otherwise then discard",
"set groups top interfaces lo0 unit 0 family inet filter input protect_re",
```

Verify Protection of Routing Engine Configuration

IN THIS SECTION

Protection of Routing Engine (Trusted Networks Configuration)

Configuration commands (CLI)

```
set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-list
10-216-192-1_32
set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-list
100-100-100-2_32
set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-list
8-8-8-8_32
set groups top firewall family inet filter protect_re term trusted_hosts then accept
set groups top firewall family inet filter protect_re term otherwise then log
set groups top firewall family inet filter protect_re term otherwise then syslog
set groups top firewall family inet filter protect_re term otherwise then discard
```

APIs

```
“switch_mgmt”:  {
    “protect_re”: {
        “enabled”: true,
        “trusted_hosts”: [
            “10.216.192.1”,
            “100.100.100.2”,
            “8.8.8.8”
        ],
        “allowed_services”: [],
        “custom”: []
    },
```

Use the `show bgp summary` command to get a summary of the status of BGP connections:

```
{master:0}
mist@Border-switch-R2-U21> show bgp summary
```

```
Warning: License key missing; One or more members of the VC require ‘bgp’ license
Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
10          6          0          0          0          0
bgp.evpn.0
68         34          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
10.255.240.3          65002        101        103       0       1       42:08 Establ
inet.0: 3/5/5/0
10.255.240.5          65003         35         33       0       3       11:51 Establ
inet.0: 3/5/5/0
100.100.100.2         65002        206        209       0       0     1:06:18 Establ
bgp.evpn.0: 25/34/34/0
default-switch.evpn.0: 22/30/30/0
default_evpn.evpn.0: 0/0/0/0
100.100.100.3         65003         57         55       0       3       11:48 Establ
bgp.evpn.0: 9/34/34/0
default-switch.evpn.0: 8/30/30/0
default_evpn.evpn.0: 0/0/0/0
```

To test the Trusted Networks functionality, ping 100.100.100.2 from the switch, as shown below. You can see that all the transmitted packets are received without any packet loss.

```
mist@Border-switch-R2-U21> ping 100.100.100.2
PING 100.100.100.2 (100.100.100.2): 56 data bytes
64 bytes from 100.100.100.2: icmp_seq=0 ttl=64 time=2.695 ms
64 bytes from 100.100.100.2: icmp_seq=1 ttl=64 time=8.756 ms
64 bytes from 100.100.100.2: icmp_seq=2 ttl=64 time=13.312 ms
64 bytes from 100.100.100.2: icmp_seq=3 ttl=64 time=9.025 ms

--- 100.100.100.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 2.695/8.447/13.312/3.781 ms

{master:0}
mist@Border-switch-R2-U21> ssh root@100.100.100.3

{master:0}
```

```
mist@Border-switch-R2-U21> ssh root@100.100.100.2
Password:
Last login: Fri Feb  3 04:57:20 2023 from 10.255.240.2
--- JUNOS 21.3R1.9 Kernel 64-bit  JNPR-12.1-20210828.6e5b1bf_buil
root@CORE-1:RE:0%
```

Also, ping or ssh a network other than the trusted networks. As you can see below, the ping shows 100 percent packet loss.

```
mist@Border-switch-R2-U21> ssh root@100.100.100.3

{master:0}
mist@Border-switch-R2-U21> ssh root@100.100.100.4

{master:0}
mist@Border-switch-R2-U21> ping 100.100.100.3
PING 100.100.100.3 (100.100.100.3): 56 data bytes

--- 100.100.100.3 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss

{master:0}
mist@Border-switch-R2-U21> ping 100.100.100.4
PING 100.100.100.4 (100.100.100.4): 56 data bytes

--- 100.100.100.4 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

Protection of Routing Engine (Trusted Services Configuration)

Configuration commands (CLI)

```
"set groups top interfaces lo0 unit 0 family inet filter input protect_re",
"set groups top firewall family inet filter protect_re term allow_mist_obssh from source-port
[ 2200 ]",
"set groups top firewall family inet filter protect_re term allow_mist_obssh then accept",
"set groups top firewall family inet filter protect_re term allow_dhcp from source-port [ 67
68 ]",
"set groups top firewall family inet filter protect_re term allow_dhcp from destination-port
[ 67 68 ]",
"set groups top firewall family inet filter protect_re term allow_dhcp from protocol udp",
"set groups top firewall family inet filter protect_re term allow_dhcp then accept",
"set groups top firewall family inet filter protect_re term allow_bgp from source-prefix-list
bgp_neighbors",
"set groups top firewall family inet filter protect_re term allow_bgp from source-prefix-list
bgp_vrf_neighbors",
"set groups top firewall family inet filter protect_re term allow_bgp from destination-port
[ 179 ]",
"set groups top firewall family inet filter protect_re term allow_bgp from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_bgp then accept",
"set groups top firewall family inet filter protect_re term allow_bfd from source-prefix-list
bgp_neighbors",
"set groups top firewall family inet filter protect_re term allow_bfd from source-prefix-list
bgp_vrf_neighbors",
"set groups top firewall family inet filter protect_re term allow_bfd from destination-port
[ 3784 4784 ]",
"set groups top firewall family inet filter protect_re term allow_bfd from protocol udp",
"set groups top firewall family inet filter protect_re term allow_bfd then accept",
"set firewall family inet filter protect_re term allow_ntp_src from protocol udp",
"set firewall family inet filter protect_re term allow_ntp_src from source-port 123",
"set firewall family inet filter protect_re term allow_ntp_src then accept",
"set firewall family inet filter protect_re term allow_ntp_dest from protocol udp",
"set firewall family inet filter protect_re term allow_ntp_dest from destination-port 123",
"set firewall family inet filter protect_re term allow_ntp_dest then accept",
"set groups top firewall family inet filter protect_re term allow_dns from source-port [ 53 ]",
"set groups top firewall family inet filter protect_re term allow_dns from protocol [ tcp
udp ]", "set groups top firewall family inet filter protect_re term allow_dns then accept",
"set groups top firewall family inet filter protect_re term allow_radius from source-prefix-list
radius_servers",
"set groups top firewall family inet filter protect_re term allow_radius from destination-port
[ 1812 1813 ]",
"set groups top firewall family inet filter protect_re term allow_radius from protocol udp",
"set groups top firewall family inet filter protect_re term allow_radius then accept",
"set groups top firewall family inet filter protect_re term allow_tacacs from source-prefix-list
tacacs_servers",
"set groups top firewall family inet filter protect_re term allow_tacacs from destination-port
[ 49 ]",
"set groups top firewall family inet filter protect_re term allow_tacacs from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_tacacs then accept",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from source-
prefix-list snmp_clients",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from destination-
port [ 161 10161 ]",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from protocol udp",
"set groups top firewall family inet filter protect_re term allow_snmp_clients then accept",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-list 10-216-192-1_32",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-list 100-100-100-2_32",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-list 8-8-8-8_32",
"set groups top firewall family inet filter protect_re term trusted_hosts then accept",
"set groups top firewall family inet filter protect_re term allow_ssh from destination-port [ 22 ]",
"set groups top firewall family inet filter protect_re term allow_ssh from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_ssh then accept",
"set groups top firewall family inet filter protect_re term allow_icmp from protocol icmp",
"set groups top firewall family inet filter protect_re term allow_icmp then accept",
"set groups top firewall family inet filter protect_re term otherwise then discard",
```

APIs

```


“switch_mgmt”:  {
    “protect_re”: {
        “enabled”: true,
        “trusted_hosts”: [
            “10.216.192.1”,
            “100.100.100.2”,
            “8.8.8.8”
        ],
        “allowed_services”: [
            “ssh”,
            “icmp”
        ],
        “custom”: []
    },
```

To test the trusted services configuration, log in to a device which is not on the trusted network.

```
mist@Distribution-2-R2-U07-> ping 100.100.100.1
PING 100.100.100.1 (100.100.100.1): 56 data bytes
64 bytes from 100.100.100.1: icmp_seq=0 ttl=63 time=36.941 ms
```

```
64 bytes from 100.100.100.1: icmp_seq=1 ttl=63 time=45.158 ms

--- 100.100.100.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 36.941/41.050/45.158/4.108 ms

{master:0}
mist@Distribution-2-R2-U07-> ssh root@100.100.100.1
Password:
Last login: Fri Feb  3 07:23:35 2023 from 10.216.201.35
--- JUNOS 22.2R1.12 Kernel 64-bit  JNPR-12.1-20220623.dbb31e0_buil
root@Border-switch-R2-U21:RE:0%
```

To check the discarded packet, run the following additional CLI commands on the device:

```
set groups top firewall family inet filter protect_re term otherwise then log
set groups top firewall family inet filter protect_re term otherwise then syslog

mist@Distribution-1-R2-U06-> show firewall log
Log :
Time      Filter    Action Interface           Protocol        Src Addr
Dest Addr
13:20:01  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:19:56  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:19:51  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:19:45  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:19:40  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:19:35  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:19:30  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:18:26  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:18:19  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:18:18  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
```

```
13:18:14  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:18:12  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
13:18:09  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:18:04  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
13:18:01  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
13:18:00  pfe       D      vtep.32769          UDP             0.0.0.0
255.255.255.255
13:18:00  pfe       D      vtep.32769          UDP             0.0.0.0
255.255.255.255
13:18:00  pfe       D      vtep.32769          UDP             0.0.0.0
255.255.255.255
13:18:00  pfe       D      vtep.32769          UDP             0.0.0.0
255.255.255.255
13:18:00  pfe       D      vtep.32769          UDP             0.0.0.0
255.255.255.255
13:18:00  pfe       D      vtep.32769          UDP             0.0.0.0
255.255.255.255
13:18:00  pfe       D      vtep.32769          UDP             0.0.0.0
255.255.255.255
13:18:00  pfe       D      vtep.32769          UDP             0.0.0.0
255.255.255.255
14:17:31  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:17:30  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:17:28  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:17:28  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:17:26  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:17:23  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:17:18  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:17:16  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:17:15  protect_re D     vme.0               UDP             8.8.8.8
```

```
10.216.202.6
14:17:12  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:17:10  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:17:09  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:17:07  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:17:06  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:17:05  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:17:03  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:17:02  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:17:01  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:16:57  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:16:52  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:16:51  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:16:50  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:16:46  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:16:45  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:16:44  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:16:41  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:16:41  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:16:40  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:16:38  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:16:36  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
```

```
14:16:36  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:16:31  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:16:26  protect_re D     vme.0               UDP             10.216.199.80
255.255.255.255
14:16:26  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:16:25  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:16:20  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
14:16:19  protect_re D     vme.0               UDP             8.8.8.8
10.216.202.6
14:16:16  protect_re D     vme.0               UDP             66.129.233.81
10.216.202.6
```

Read also: Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources and Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods.

QoS Configuration

SUMMARY

Understand Quality of Service (Qos) concepts and codes, and follow these steps to enable QoS in a port profile.

IN THIS SECTION

NOTE: QoS configuration is part of the switch configuration workflow described in "Configure Switches Using Templates" on page 21. This topic provides more detailed information focused solely on QoS concept and configuration steps.

Quality of Service (QoS) is a traffic-control mechanism that helps you prioritize latency-sensitive traffic (such as voice) over other traffic in a congested network. Juniper Mist enables QoS on a per interface basis. The QoS implementation generally involves the following aspects:

- Classifying traffic.
- Defining traffic-to-queue mappings (forwarding classes).
- Defining scheduler and re-write rules for each queue. These rules govern the prioritization, bandwidth control, and congestion management of the traffic on each interface.
- Applying QoS components to the interfaces.

In Juniper Mist, QoS utilizes the Behavior Aggregate (BA) classification, where the DiffServ code point (DSCP) or class of service (CoS) values in the incoming traffic govern the classification. The BA classifier maps a CoS value in the packet header to a forwarding class and loss priority.

Enabling QoS on an interface adds DSCP markings to that port based on the class and rewrite rules. The QoS mechanism maps the incoming packets with a DSCP marking to one of the seven forwarding classes listed in the following table:

Code Point/Loss Priority	Forwarding Class	Transmit Queue	Buffer Size(%)	Transmit Rate(%)	Priority
be	default-app	0	Remainder	Remainder	Low
af41/Low af42/ High af43/High cs4/Low	video	1	8	8	Low
af31/Low af32/ High af33/High cs3/Low	bizapp-af3	2	10	10	Low
af21/Low af22/ High af23/High	bizapp-af2	3	10	10	Low
af11/Low af12/ High af13/High	net-tools	4	3	3	Low
cs5/Low ef/Low	voice	7	10	10	Strict-high
nc1/Low nc2/Low	net-control	5	3	3	Low

As shown in the above table, the packet classification assigns an incoming packet to an output queue based on the packet's forwarding class. In case of traffic congestion on the link, Juniper Mist prioritizes the latency-sensitive traffic (for example, voice traffic) over other traffic (provided that the incoming traffic is marked appropriately). Juniper Mist also configures re-write rules automatically to retain markings as the packets exit the switch.

Enable QoS on a Switch Port

Enabling QoS helps you prioritize latency-sensitive traffic (such as voice) over other traffic in a congested network. You can configure QoS on a switch port from the Port Profile tile on the switch details page or switch template.

NOTE: Ensure that you enable QoS on both downstream and upstream port profiles, to obtain optimum results.

To enable QoS on a switch port:

1. To configure QoS at the organization level, click **Organization** > **Switch Templates** > *template name*. Or, if you want to configure QoS at the switch level, click **Switches** > *switch name*.

2. From the Port Profile tile, select the port profile you want to update. Or if you want to create a new port profile, click **Add Profile**.

3. In the configuration, remember to select the **QoS** check box.

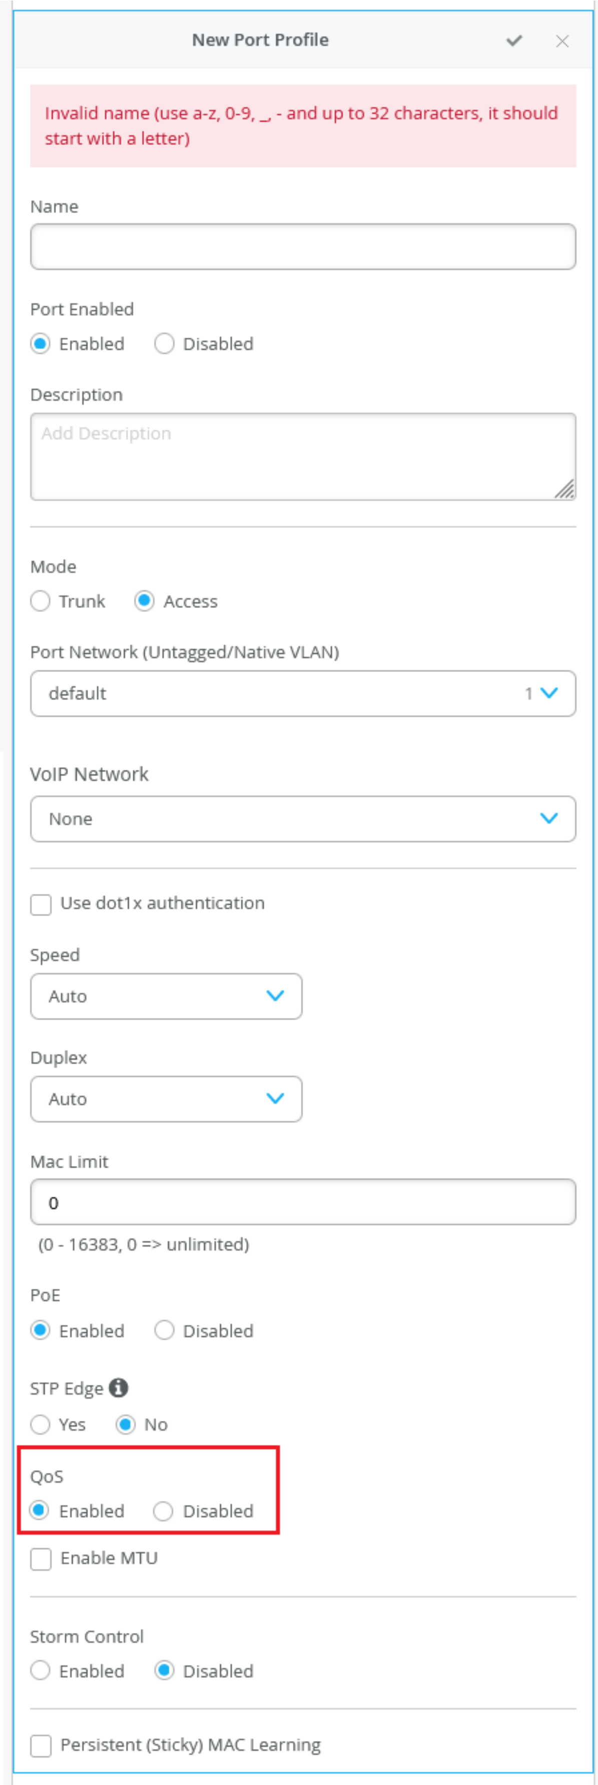

New Port Profile

Invalid name (use a-z, 0-9, _, - and up to 32 characters, it should start with a letter)

Name

Port Enabled
● Enabled ○ Disabled

Description
Add Description

Mode
○ Trunk ● Access

Port Network (Untagged/Native VLAN)
default 1

VoIP Network
None

□ Use dot1x authentication

Speed
Auto

Duplex
Auto

Mac Limit
0
(0 - 16383, 0 => unlimited)

PoE
● Enabled ○ Disabled

STP Edge
○ Yes ● No

QoS
● Enabled ○ Disabled

□ Enable MTU

Storm Control
○ Enabled ● Disabled

□ Persistent (Sticky) MAC Learning

4. Save the configuration by clicking the tick mark on the upper right of the port profile configuration window.

5. After configuring QoS in the port profile, assign the profile to the switch port on which you want to configure QoS. You can do that from the Port Config tab in the Select Switches section of a switch configuration template (See "Create a Switch Configuration Template" on page 22) or from the Port Configuration section on the Switch Details page ("Switch Details" on page 294).

Override QoS

You also have the option to override the QoS configuration on the WLAN settings page (**Site** > **WLANs** > ***WLAN name***). To override the QoS configuration, select the **Override QoS** check box and choose a wireless access class (see WLAN Options). The downstream traffic (AP > client) gets marked with the override access class value specified. The override configuration doesn't support upstream traffic (client > AP).

For further details on QoS on Juniper EX Switches, please see Example: Configuring CoS on EX Series Switches. If required, any additional QoS configuration updates can be done via CLIs in the Additional CLI Commands section in the switch details page.

Verify QoS Settings (API)

The following example has `“enable_qos”: true` set for the port profiles `qos-test` and `uplink`. This indicates that the port profile has QoS enabled.

```
"port_usages": {
        "qos-test": {
            "name": "qos-test",
            "mode": "access",
            "disabled": false,
            "port_network": "vl10",
            "voip_network": null,
            "stp_edge": false,
            "all_networks": false,
            "networks": [],
            "port_auth": null,
            "speed": "auto",
            "duplex": "auto",
            "mac_limit": 0,
            "poe_disabled": false,
            "enable_qos": true
        },
        "uplink": {
            "mode": "trunk",
            "all_networks": true,
            "stp_edge": false,
            "port_network": "vlan3",
            "voip_network": null,
            "name": "uplink",
            "disabled": false,
            "networks": [],
            "port_auth": null,
            "speed": "auto",
            "duplex": "auto",
            "mac_limit": 0,
            "poe_disabled": false,
            "enable_qos": true
        }
    },
```

Verify QoS Configuration Through the CLI

The following is a sample QoS configuration on a switch:

```
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class bizapp-af2 loss-priority high code-points af22
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class bizapp-af2 loss-priority high code-points af23
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class bizapp-af2 loss-priority low code-points af21
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class bizapp-af3 loss-priority high code-points af32
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class bizapp-af3 loss-priority high code-points af33
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class bizapp-af3 loss-priority low code-points af31
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class bizapp-af3 loss-priority low code-points cs3
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class default-app loss-priority low code-points be
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class net-control loss-priority low code-points nc1
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class net-control loss-priority low code-points nc2
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class net-tools loss-priority high code-points af12
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class net-tools loss-priority high code-points af13
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class net-tools loss-priority low code-points af11
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class video loss-priority high code-points af42
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class video loss-priority high code-points af43
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class video loss-priority low code-points af41
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class video loss-priority low code-points cs4
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class voice loss-priority low code-points cs5
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-class voice loss-priority low code-points ef
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default import default
set groups mist-qos-default class-of-service forwarding-classes queue 0 default-app
set groups mist-qos-default class-of-service forwarding-classes queue 1 video
set groups mist-qos-default class-of-service forwarding-classes queue 2 bizapp-af3
set groups mist-qos-default class-of-service forwarding-classes queue 3 bizapp-af2
set groups mist-qos-default class-of-service forwarding-classes queue 4 net-tools
set groups mist-qos-default class-of-service forwarding-classes queue 5 voice
set groups mist-qos-default class-of-service forwarding-classes queue 7 net-control
set groups mist-qos-default class-of-service interfaces ge-0/0/0 scheduler-map sched-maps-default
set groups mist-qos-default class-of-service interfaces ge-0/0/0 unit 0 classifiers dscp dscp-classifier-default
set groups mist-qos-default class-of-service interfaces ge-0/0/0 unit 0 rewrite-rules dscp dscp-rewriter-default
set groups mist-qos-default class-of-service interfaces ge-0/0/9 scheduler-map sched-maps-default
set groups mist-qos-default class-of-service interfaces ge-0/0/9 unit 0 classifiers dscp dscp-classifier-default
set groups mist-qos-default class-of-service interfaces ge-0/0/9 unit 0 rewrite-rules dscp dscp-rewriter-default
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewrite-default import default
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class bizapp-af2 loss-priority low code-point af21
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class bizapp-af3 loss-priority low code-point af31
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class default-app loss-priority low code-point be
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class net-control loss-priority low code-point nc1
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class net-tools loss-priority low code-point af11
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class video loss-priority low code-point af41
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class voice loss-priority low code-point ef
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class bizapp-af2 scheduler bizapp-af2-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class bizapp-af3 scheduler bizapp-af3-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class default-app scheduler default-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class net-control scheduler net-control-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class net-tools scheduler net-tools-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class video scheduler video-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class voice scheduler voice-scheduler
set groups mist-qos-default class-of-service schedulers bizapp-af2-scheduler buffer-size percent 10
set groups mist-qos-default class-of-service schedulers bizapp-af2-scheduler priority low
set groups mist-qos-default class-of-service schedulers bizapp-af2-scheduler transmit-rate percent 10
set groups mist-qos-default class-of-service schedulers bizapp-af3-scheduler buffer-size percent 10
set groups mist-qos-default class-of-service schedulers bizapp-af3-scheduler priority low
set groups mist-qos-default class-of-service schedulers bizapp-af3-scheduler transmit-rate percent 10
set groups mist-qos-default class-of-service schedulers default-scheduler buffer-size remainder
set groups mist-qos-default class-of-service schedulers default-scheduler priority low
set groups mist-qos-default class-of-service schedulers default-scheduler transmit-rate remainder
set groups mist-qos-default class-of-service schedulers net-control-scheduler buffer-size
percent 3
set groups mist-qos-default class-of-service schedulers net-control-scheduler priority low
set groups mist-qos-default class-of-service schedulers net-control-scheduler transmit-rate
percent 3
set groups mist-qos-default class-of-service schedulers net-tools-scheduler buffer-size percent 3
set groups mist-qos-default class-of-service schedulers net-tools-scheduler priority low
set groups mist-qos-default class-of-service schedulers net-tools-scheduler transmit-rate
percent 3
set groups mist-qos-default class-of-service schedulers video-scheduler buffer-size percent 8
set groups mist-qos-default class-of-service schedulers video-scheduler priority low
set groups mist-qos-default class-of-service schedulers video-scheduler transmit-rate percent 8
set groups mist-qos-default class-of-service schedulers voice-scheduler buffer-size percent 10
set groups mist-qos-default class-of-service schedulers voice-scheduler priority strict-high
set groups mist-qos-default class-of-service schedulers voice-scheduler shaping-rate percent 10
```

To verify the traffic-matching QoS policies and their corresponding queue counters:

1. Review the current interface statistics and CoS information by running the following command:

```
root@ex2300-home> show interfaces ge-0/0/0 extensive
......
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
    2                                0                    0                    0
    3                                0                    0                    0
    4                                0                    0                    0
    5                                0                    0                    0
    7                                0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   default-app
    1                   video
    2                   bizapp-af3
    3                   bizapp-af2
    4                   net-tools
    5                   voice
    7                   net-control
  ......
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth              Buffer Priority   Limit
                              %            bps     %           usec
    0 default-app             r              r     r              0      low    none
    1 video                   8       80000000     8              0      low    none
    2 bizapp-af3             10      100000000    10              0      low    none
    3 bizapp-af2             10      100000000    10              0      low    none
    4 net-tools               3       30000000     3              0      low    none
    5 voice                   r              r    10              0 strict-high    none
    7 net-control             3       30000000     3              0      low    none
  Interface transmit statistics: Disabled
```

2. Generate some video and voice traffic. The device marks the traffic with DSCP values (queue 1 for video traffic and queue 5 for voice traffic).

```
ping 8.8.8.8 -I eth0 -Q 184
PING 8.8.8.8 (8.8.8.8) from 10.0.0.2 eth0: 56(84) bytes of data.

53 packets transmitted, 53 received, 0% packet loss, time 140ms
rtt min/avg/max/mdev = 2.421/2.811/5.064/0.428 ms
```

```
ping 8.8.8.8 -I eth0 -Q 136
PING 8.8.8.8 (8.8.8.8) from 10.0.0.2 eth0: 56(84) bytes of data.

62 packets transmitted, 62 received, 0% packet loss, time 157ms
rtt min/avg/max/mdev = 2.396/3.103/6.578/0.609 ms
```

3. Run the `show interfaces ge-0/0/0 extensive` command again. You can view the packet counts displayed under Queued Packets and Transmitted Packets.

```
root@ex2300-home> show interfaces ge-0/0/0 extensive
.......
  Egress queues: 8 supported, 7 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                             9821                 9821                    0
```

```
   1                               62                   62                    0
   2                                0                    0                    0
   3                             7185                 7185                    0
   4                                0                    0                    0
   5                               53                   53                    0
   7                                0                    0                    0
 Queue number:         Mapped forwarding classes
   0                   default-app
   1                   video
   2                   bizapp-af3
   3                   bizapp-af2
   4                   net-tools
   5                   voice
   7                   net-control
.......
 CoS information:
   Direction : Output
   CoS transmit queue               Bandwidth               Buffer Priority   Limit
                             %            bps     %           usec
   0 default-app             r              r     r              0      low    none
   1 video                   8       80000000     8              0      low    none
   2 bizapp-af3             10      100000000    10              0      low    none
   3 bizapp-af2             10      100000000    10              0      low    none
   4 net-tools               3       30000000     3              0      low    none
   5 voice                   r              r    10              0 strict-high    none
   7 net-control             3       30000000     3              0      low    none
 Interface transmit statistics: Disabled
```

See also: Example: Configuring CoS on EX Series Switches

Configure LAG with Wired Assurance

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a *link aggregation group (LAG)* or *bundle*.

Aggregating multiple links between physical interfaces creates a single logical point-to-point trunk link or a LAG. The LAG balances traffic across the member links within an aggregated Ethernet (ae) bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

To configure LAG on a switch via Mist:

1. Navigate to the switch configuration where you want to configure LAG.

   To configure at the organization level, go to **Organization** > **Wired** > **Switch Templates** > *Template Name*.

   To configure at the site level, go to **Site** > **Wired** > **Switch Configuration** > *Template Name*.

   To configure LAG on individual switch (from the switch details page), click **Switches** > *Switch Name*.
2. (Applicable only to the organization-level or site-level configuration) On the switch template, navigate to the Select Switches Configuration section and then click **Add Rule**.

   a. Go to the Select Switches Configuration section and then click **Add Rule**.

      If you want to modify an existing switch rule, click the rule from he left menu of the Select Switches Configuration section.

   b. Navigate to the **Port Config** tab and then click **Add Port Configuration**.

      The New Port Configuration window appears.
3. (Applicable only to the switch-level configuration) If you are configuring LAG on an individual switch, navigate to the **Port** tile on the switch details page and then click **Add Port Configuration**.

   The New Port Configuration window appears.
4. In the **Port IDs** field on the New Port Configuration window, specify the interface names (for example, ge-0/0/1, ge-0/0/4) that you want to include in the port aggregation or LAG configuration.

   > **NOTE**: The number of interfaces that you can group into a LAG and the total number of LAGs that a switch supports vary depending on switch model.

5. Select a port profile from the **Configuration Profile** drop-down list.

   > **NOTE**:
   >
   > - The Q-in-Q profile does not support link aggregation.
   > - A port does not support port aggregation if the port profile associated with it has the Dot1x option enabled.

6. Select **Enabled** under Port Aggregation.
7. Specify an AE index for the aggregated port.

   > **NOTE**:

- Ensure that the AE index does not overlap across different ports between device, site and organization level template, and campus fabric configuration.

- When a device is part of a campus fabric, the 'esi-lag' option appears in the Port Aggregation section. Enable this option on both switches in the pair, using the same AE index. Doing so configures a common ESI on both switches. For more information, refer to Ethernet Segment Identifiers, ESI Types, and LACP in EVPN LAGs.

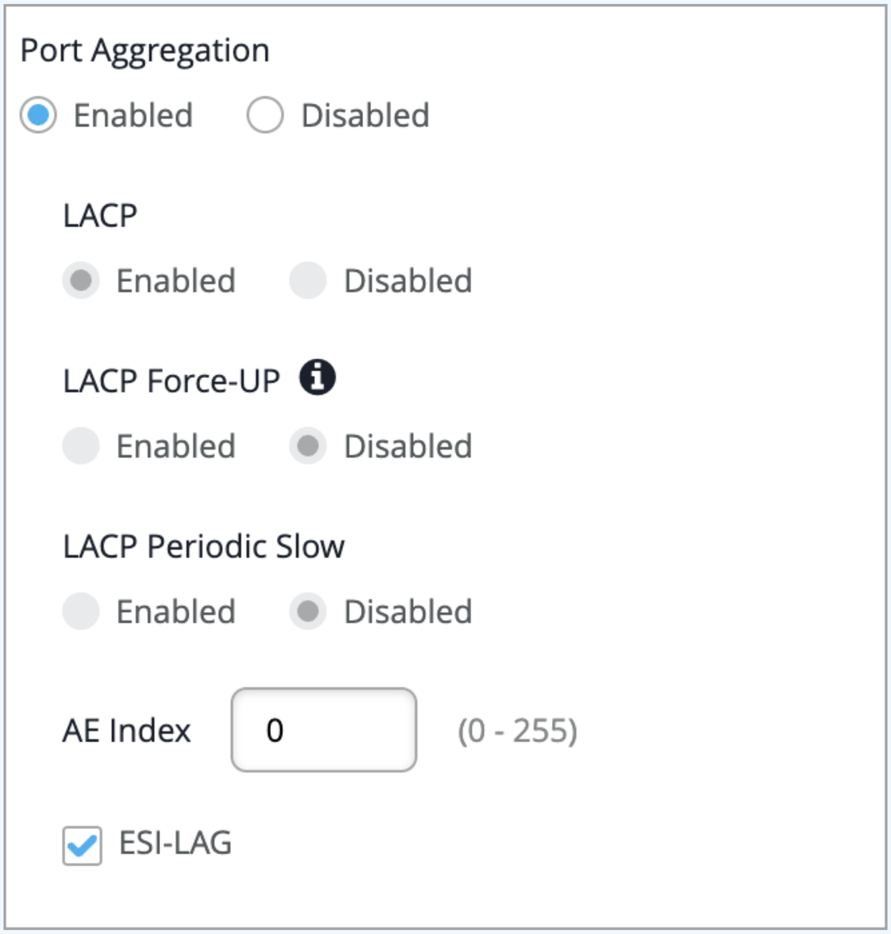

Port Aggregation

● Enabled ○ Disabled

LACP

● Enabled ○ Disabled

LACP Force-UP ℹ

○ Enabled ● Disabled

LACP Periodic Slow

○ Enabled ● Disabled

AE Index 0 (0 - 255)

☑ ESI-LAG

8. Configure other settings such as the following:

   - LACP. You can use LAG with or without LACP enabled. If the device on the other end doesn't support LACP, you can disable LACP here.

   - The LACP force-up state for the switch. This configuration sets the state of the interface as up when the peer has limited LACP capability.

   - An LACP packet transmission interval. If you configure the LACP Periodic Slow option on an AE interface, the LACP packets are transmitted every 30 seconds. By default, the interval is set to fast in which the packets are transmitted every second.

9. Click the check mark on the upper right of the New Port Configuration window to save the port configuration.
10. If you are configuring LAG at the template level, click the check mark on the right side of the switch rules section.
11. Save the configuration.
12. To verify the configuration:

    a. Navigate to the switch details page by clicking **Switches** > *Switch Name*.

    b. On the Front Panel tab on the switch details page, hover over the ports on which you configured port aggregation.

    The ae ports will appear with a purple halo around it.

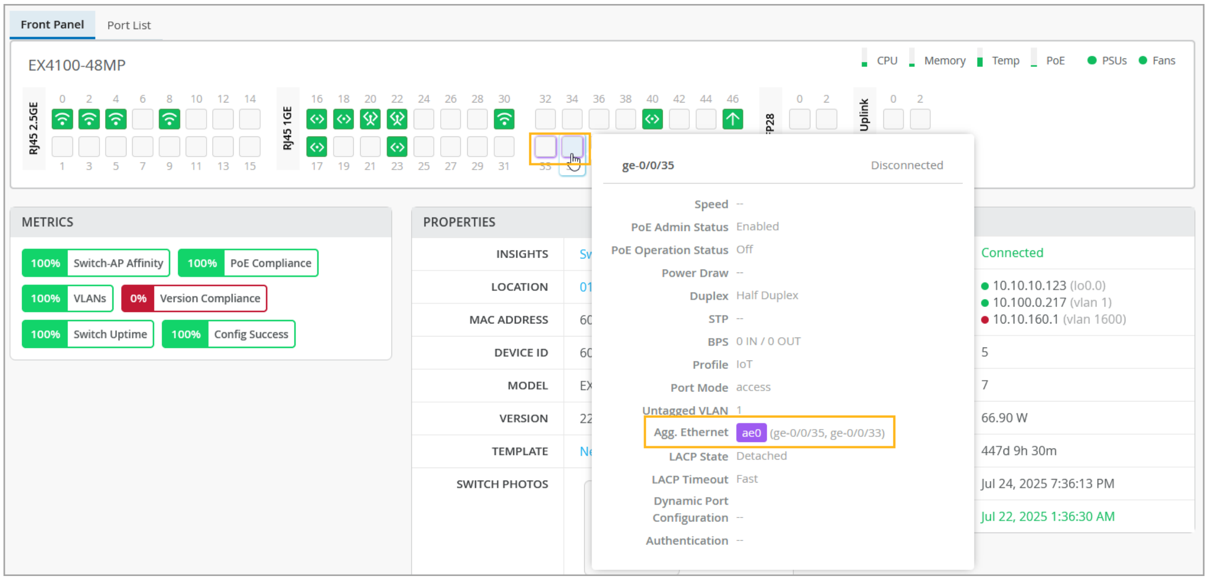

100% Switch-AP Affinity
100% PoE Compliance
100% VLANs
0% Version Compliance
100% Switch Uptime
100% Config Success

INSIGHTS Sw
LOCATION 01
MAC ADDRESS 60
DEVICE ID 60
MODEL EX
VERSION 22
TEMPLATE Ne
SWITCH PHOTOS

ge-0/0/35 Disconnected

Speed --
PoE Admin Status Enabled
PoE Operation Status Off
Power Draw --
Duplex Half Duplex
STP --
BPS 0 IN / 0 OUT
Profile IoT
Port Mode access
Untagged VLAN 1
Agg. Ethernet ae0 (ge-0/0/35, ge-0/0/33)
LACP State Detached
LACP Timeout Fast
Dynamic Port Configuration --
Authentication --

Connected
● 10.10.10.123 (lo0.0)
● 10.100.0.217 (vlan 1)
● 10.10.160.1 (vlan 1600)
5
7
66.90 W
447d 9h 30m
Jul 24, 2025 7:36:13 PM
Jul 22, 2025 1:36:30 AM

Configure SNMP on Switches

SUMMARY

Follow these steps to configure SNMP in switch templates, the site configuration, or the device settings.

IN THIS SECTION

NOTE: SNMP configuration is part of the switch configuration workflow described in "Configure Switches Using Templates" on page 21. This topic provides more detailed information focused solely on SNMP configuration steps.

Simple Network Management Protocol (SNMP) enables network administrators to monitor and manage network-connected devices in IP networks (see SNMP Architecture and SNMP MIBs Overview).

In a switch, SNMP is disabled by default. If required, you can enable it at the organization level (from the organization level switch templates), site level (from the site level switch templates) at device level (from the switch details page.)

Mist allows you to configure SNMP V2 or SNMP V3 settings.

Configure SNMP at the Organization Level

If you want to apply SNMP to all the switches in the entire organization (all sites), you can do that at the organization level.

To configure SNMP at the organization level:

1. Click **Organization** > **Switch Templates**.
2. Click **Create Template** if you want to configure SNMP as part of a new switch template. For more information, see "Create a Switch Configuration Template" on page 22.

   Or, if you want to update an existing template with SNMP configuration, click that template to open it.
3. On the SNMP tile (in the All Switches Configuration section), click the **Enabled** check box.

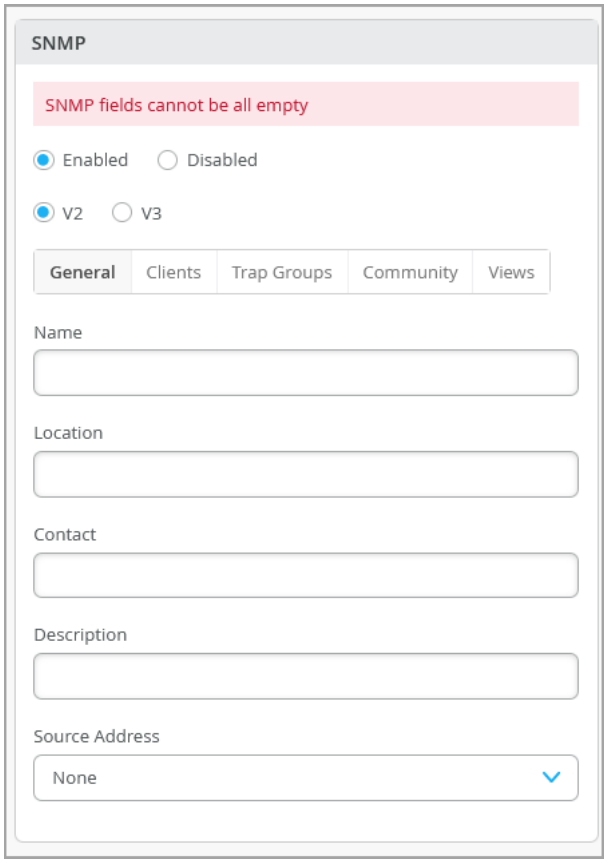

SNMP

SNMP fields cannot be all empty

● Enabled ○ Disabled

● V2 ○ V3

General | Clients | Trap Groups | Community | Views

Name

Location

Contact

Description

Source Address

None

4. Specify the SNMP information, which includes General information, Clients, Trap Groups, Community, and Views. Mist supports the SNMP versions V2 and V3. For the field descriptions, refer to the SNMP information in the **All Switches Configuration Options** table in "All Switches" on page 30.

   The configuration fields appear based on the selection of SNMP version (V2 or V3).
5. Save the template.
6. Assign the template to a site. For instructions to do so, see "Assign a Template to Sites" on page 22.

Configure SNMP at the Site Level

You can choose to have SNMP configuration to specific sites in an organization. If the SNMP is disabled at the organization level but you want to enable it at a particular site, you can override the organization templates settings inherited by that site. If you configure SNMP at the site level, the configuration gets applied to all the switches in that site. Any update in the site-level settings does not override the values in the associated organization template. The change is applied only to the selected site.

To configure SNMP at the site level:

1. Click **Site** > **Switch Configuration**.

2. Click the site, where you want to configure SNMP, to open it.
3. Scroll down to the SNMP tile in the All Switches Configuration section.
4. Select the **Override Configuration Template** check box.
5. Click the **Enabled** check box in the SNMP tile.

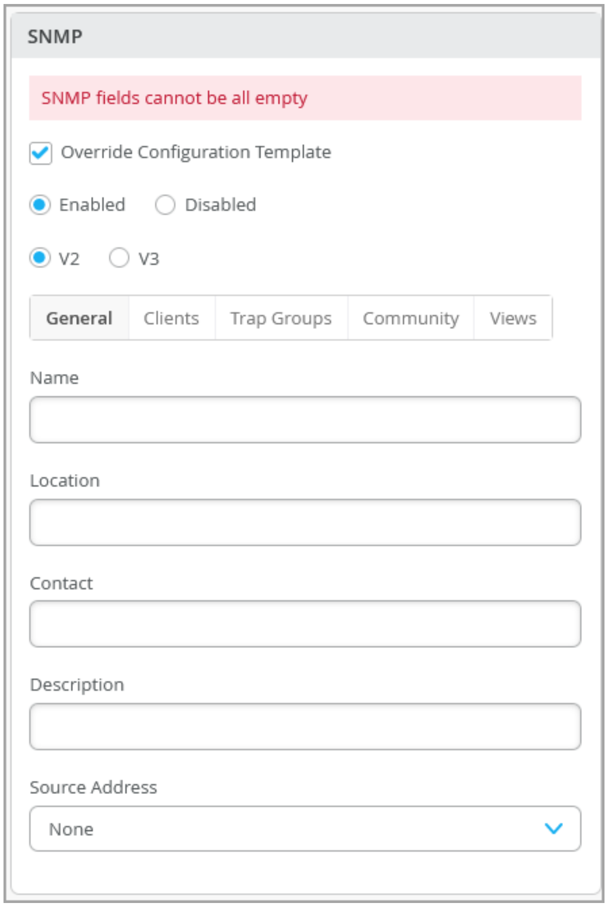

SNMP

SNMP fields cannot be all empty

☑ Override Configuration Template

● Enabled ○ Disabled

● V2 ○ V3

Name

Location

Contact

Description

Source Address

None

6. Specify the SNMP information, which includes General information, Clients, Trap Groups, Community, and Views. Mist supports the SNMP versions V2 and V3. For the field descriptions, refer to the SNMP information in the **All Switches Configuration Options** table in "All Switches" on page 30.

   The configuration fields appear based on the selection of SNMP version (V2 or V3).
7. Save the configuration.

Configure SNMP at the Device Level

If you want to configure SNMP for specific switches in a site, you can do that from the switch details page. If the SNMP is disabled at the site level but you want to enable it for a specific switch, you can override the site templates settings inherited by that switch. Any update in the switch-level settings does not override the values in the associated site or organization template settings. The change is applied only to the selected switch.

Mist allows you to configure a portion of SNMP settings at the switch level and use the remaining portion from the template. This means you can merge the SNMP values configured at the switch-level with the SNMP values that the switch inherited from a site-level or organization-level template. This feature is helpful when you want the switch to use some SNMP values from the switch-level configuration and some from the associated template. For example, you can use the name and location from the switch-level configuration and everything else from the associated template.

NOTE: Before you configure SNMP at the switch level, ensure that the switch is in connected state with stable connectivity to cloud.

To configure SNMP at the switch level:

1. Click **Switches** > ***Switch Name***.
2. Click the switch (to be updated) from the list to open it.
3. Scroll down to the SNMP tile in the Services section.
4. Select the **Override Site/Template Settings** check box.
5. Click the **Enabled** check box in the SNMP tile.

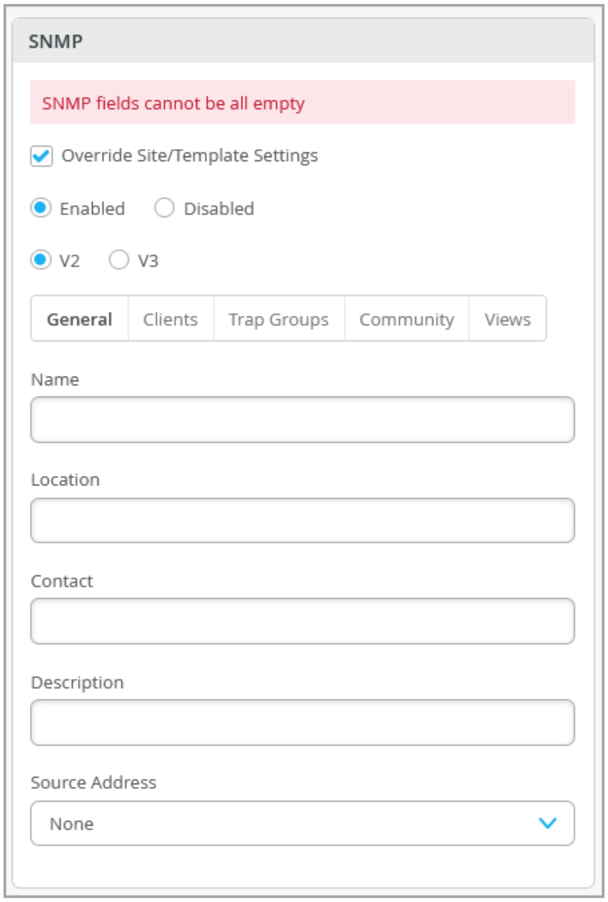

SNMP

SNMP fields cannot be all empty

☑ Override Site/Template Settings

● Enabled ○ Disabled

● V2 ○ V3

Name

Location

Contact

Description

Source Address

None

6. Specify the SNMP information, which includes General information, Clients, Trap Groups, Community, and Views. Mist supports the SNMP versions V2 and V3. For the field descriptions, refer to the SNMP information in the **All Switches Configuration Options** table in "All Switches" on page 30.

   The configuration fields appear based on the selection of SNMP version (V2 or V3).

7. Save the configuration.

Configure DHCP Server or Relay on a Switch

SUMMARY

Follow these steps to configure DHCP server or relay on a switch.

IN THIS SECTION

Prerequisites

Before configuring the DHCP server or relay, ensure the following:

- The VLAN for which DHCP server will be configured on switch is assigned to the ports connecting to the DHCP clients. You can do this by applying a relevant port profile to the port. For more information about port profiles, see "Port Profiles" on page 71.
- The switch has a **Static** IP Configuration or Additional IP configuration for the network to run DHCP server. Or the L3 Interface is configured for DHCP server or relay.

See "Wired Clients" on page 309 for information on displaying DHCP client properties in the Wired Clients list.

Configure DHCP Server

To configure DHCP server on a switch:

1. Navigate to **Switches** on the left menu and then click a Switch from the list.
   The switch details page is displayed.
2. On the switch details page, scroll down to the **DHCP Server / Relay** tile in the Services section.

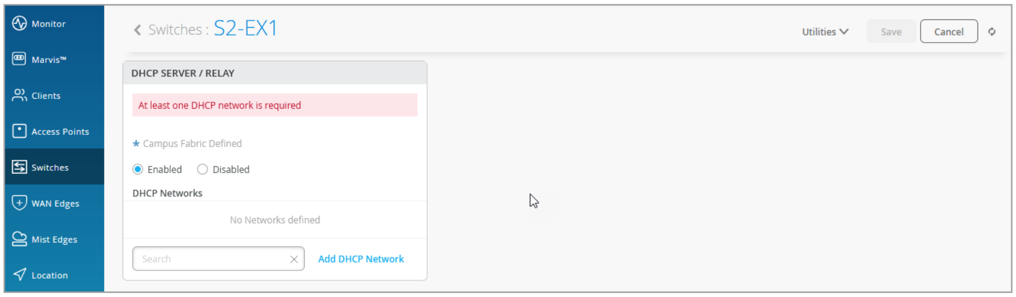

At least one DHCP network is required

* Campus Fabric Defined

● Enabled ○ Disabled

DHCP Networks

No Networks defined

3. On the **DHCP Server / Relay** tile, select the **Enabled** check box and then click **Add DHCP Network**.
   The DHCP Server / Relay configuration window is displayed.
4. Ensure that **Server** is selected as configuration **Type**.
5. From the Network drop-down list, select a network for the DHCP server.
6. Specify the following fields:

- IP Start—The starting IP address within the DHCP IP address assignment pool.
- IP End—The ending IP address within the DHCP IP address assignment pool.
- Gateway—Default gateway for the DHCP client. Usually, it is the switch IRB IP address for the corresponding network.
- DNS Servers (Optional)—You can add up to three DNS IP addresses in a comma separated format.
- DNS Suffix (Optional)—You can add up to three domain suffixes in a comma separated format.

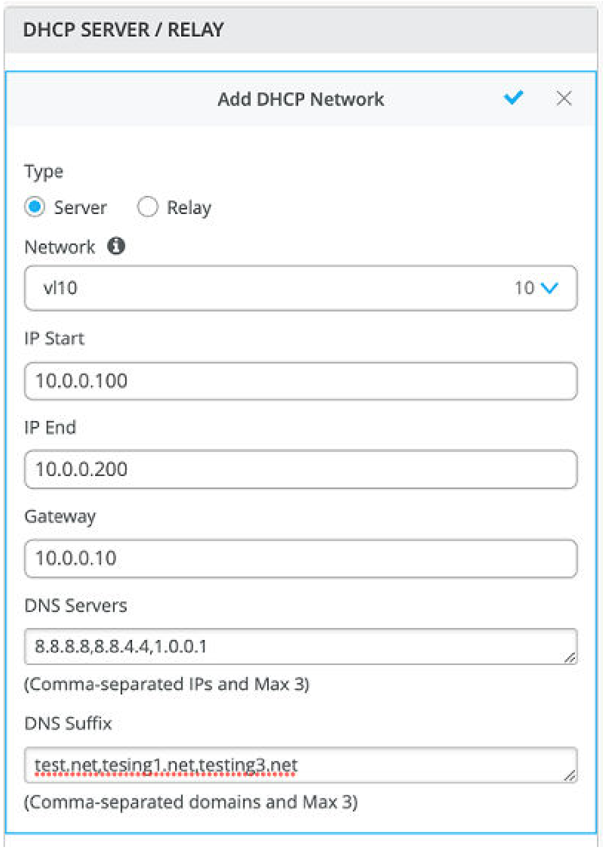

DHCP SERVER / RELAY

Add DHCP Network ✓ ×

Type

● Server ○ Relay

Network: vl10 10

IP Start: 10.0.0.100

IP End: 10.0.0.200

Gateway: 10.0.0.10

DNS Servers: 8.8.8.8,8.8.4.4,1.0.0.1

(Comma-separated IPs and Max 3)

DNS Suffix: test.net,tesing1.net,testing3.net

(Comma-separated domains and Max 3)

If you do not configure the DNS Servers and DNS Suffix, Mist auto-generates the configuration by taking DNS servers and DNS Suffix configured at the switch level.

7. Save the changes.

   The following configuration will be applied to the API `/sites/:site-id/devices/device-id`:

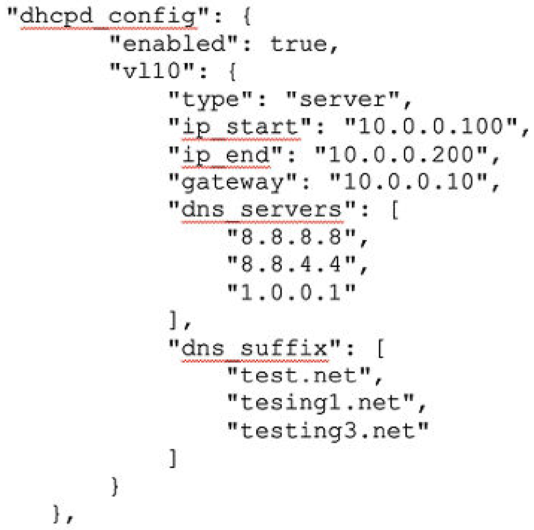

```
"dhcpd_config": {
        "enabled": true,
        "vl10": {
            "type": "server",
            "ip_start": "10.0.0.100",
            "ip_end": "10.0.0.200",
            "gateway": "10.0.0.10",
            "dns_servers": [
                "8.8.8.8",
                "8.8.4.4",
                "1.0.0.1"
            ],
            "dns_suffix": [
                "test.net",
                "tesing1.net",
                "testing3.net"
            ]
        }
    },
```

Check the corresponding intended configuration to be pushed to the following API: `/sites/:site-id/devices/device-id/config_cmd`

The DHCP server is enabled on the irb.10 interface for the vl10.

```
set groups top system services dhcp-local-server group vl10 interface irb.10
```

The address assignment pool configuration looks like the below:

```


"set groups top access address-assignment pool vl10 family inet network 10.0.0.0/24",
"set groups top access address-assignment pool vl10 family inet range vl10 low 10.0.0.100",
"set groups top access address-assignment pool vl10 family inet range vl10 high 10.0.0.200",
"set groups top access address-assignment pool vl10 family inet dhcp-attributes name-server
8.8.8.8",
"set groups top access address-assignment pool vl10 family inet dhcp-attributes name-server
8.8.4.4",
"set groups top access address-assignment pool vl10 family inet dhcp-attributes name-server
1.0.0.1",
"set groups top access address-assignment pool vl10 family inet dhcp-attributes router
10.0.0.10",
"set groups top access address-assignment pool vl10 family inet dhcp-attributes option 15
array string test.net",
"set groups top access address-assignment pool vl10 family inet dhcp-attributes option 15
array string tesing1.net",
"set groups top access address-assignment pool vl10 family inet dhcp-attributes option 15
array string testing3.net,"
```

8. Verify the configuration on the switch by using the following CLI:

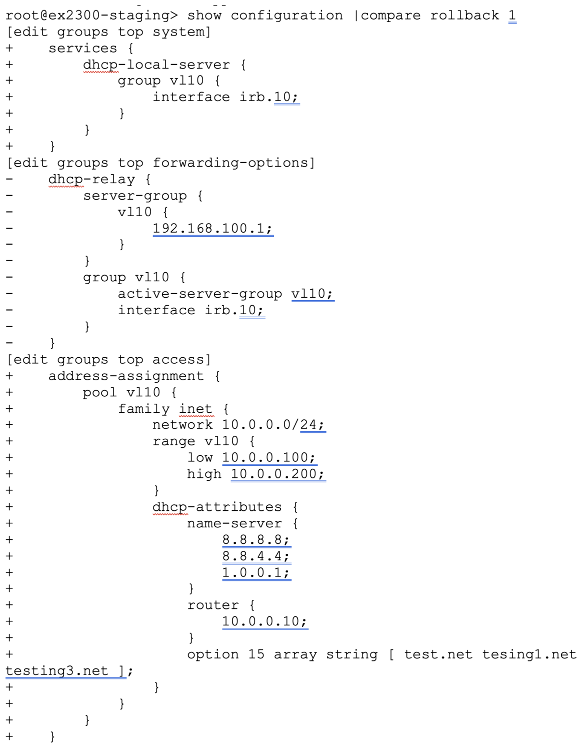

```
root@ex2300-staging> show configuration |compare rollback 1
[edit groups top system]
+    services {
+        dhcp-local-server {
+            group vl10 {
+                interface irb.10;
+            }
+        }
+    }
[edit groups top forwarding-options]
-    dhcp-relay {
-        server-group {
-            vl10 {
-                192.168.100.1;
-            }
-        }
-        group vl10 {
-            active-server-group vl10;
-            interface irb.10;
-        }
-    }
[edit groups top access]
+    address-assignment {
+        pool vl10 {
+            family inet {
+                network 10.0.0.0/24;
+                range vl10 {
+                    low 10.0.0.100;
+                    high 10.0.0.200;
+                }
+                dhcp-attributes {
+                    name-server {
+                        8.8.8.8;
+                        8.8.4.4;
+                        1.0.0.1;
+                    }
+                    router {
+                        10.0.0.10;
+                    }
+                    option 15 array string [ test.net tesing1.net
testing3.net ];
+                }
+            }
+        }
+    }
```

The DHCP clients for v110 networks should now be assigned IP address by DHCP server as shown below:

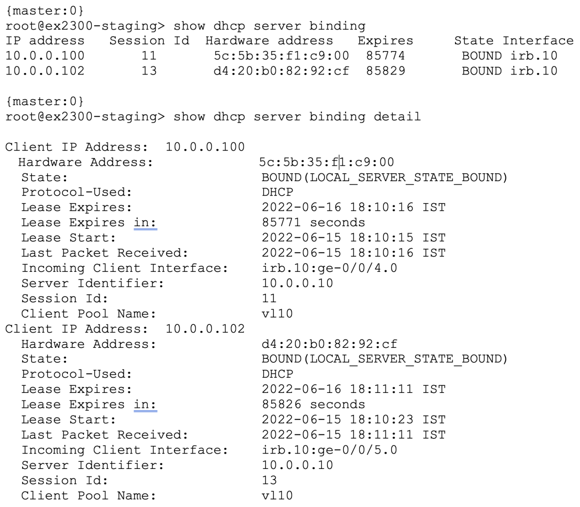

```
{master:0}
root@ex2300-staging> show dhcp server binding
IP address    Session Id  Hardware address   Expires     State Interface
10.0.0.100       11       5c:5b:35:f1:c9:00  85774       BOUND irb.10
10.0.0.102       13       d4:20:b0:82:92:cf  85829       BOUND irb.10

{master:0}
root@ex2300-staging> show dhcp server binding detail

Client IP Address:  10.0.0.100
  Hardware Address:             5c:5b:35:f1:c9:00
  State:                        BOUND(LOCAL_SERVER_STATE_BOUND)
  Protocol-Used:                DHCP
  Lease Expires:                2022-06-16 18:10:16 IST
  Lease Expires in:             85771 seconds
  Lease Start:                  2022-06-15 18:10:15 IST
  Last Packet Received:         2022-06-15 18:10:16 IST
  Incoming Client Interface:    irb.10:ge-0/0/4.0
  Server Identifier:            10.0.0.10
  Session Id:                   11
  Client Pool Name:             vl10
Client IP Address:  10.0.0.102
  Hardware Address:             d4:20:b0:82:92:cf
  State:                        BOUND(LOCAL_SERVER_STATE_BOUND)
  Protocol-Used:                DHCP
  Lease Expires:                2022-06-16 18:11:11 IST
  Lease Expires in:             85826 seconds
  Lease Start:                  2022-06-15 18:10:23 IST
  Last Packet Received:         2022-06-15 18:11:11 IST
  Incoming Client Interface:    irb.10:ge-0/0/5.0
  Server Identifier:            10.0.0.10
  Session Id:                   13
  Client Pool Name:             vl10
```

Here are some useful Junos CLI commands:

- `show dhcp server binding` or `show dhcp server binding detail`—To view the DHCP server binding information.
- `show dhcp server statistics`—To view the DHCP messages sent/received statistics.
- `clear dhcp server binding interface x-x/x/x`—To clear the DHCP client binding for a client on a particular interface.
- `clear dhcp server binding <address>`—To clear the DHCP client binding based on an IP address or MAC address of the wired client.

Configure DHCP Relay

To configure DHCP relay on a switch:

1. Navigate to **Switches** on the left menu and then click a Switch from the List.
   The switch details page is displayed.
2. On the switch details page, scroll down to the **DHCP Server / Relay** tile in the Services section.

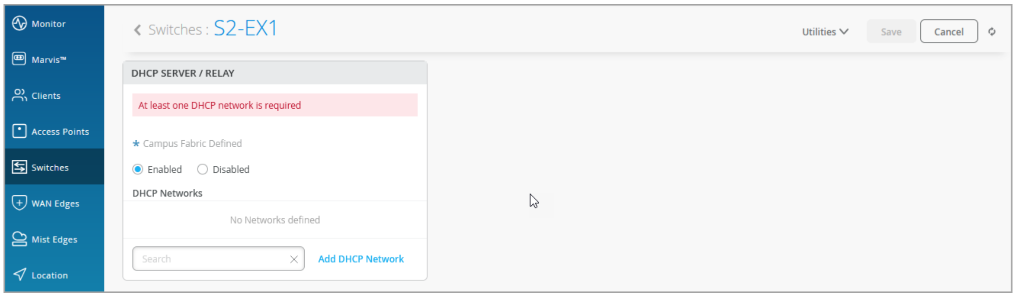

Switches : S2-EX1

DHCP SERVER / RELAY

At least one DHCP network is required

★ Campus Fabric Defined

● Enabled ○ Disabled

DHCP Networks

No Networks defined

3. On the **DHCP Server / Relay** tile, select the **Enabled** check box and then click **Add DHCP Network**.
   The DHCP Server / Relay configuration window is displayed.
4. Select **Relay** as configuration **Type**.

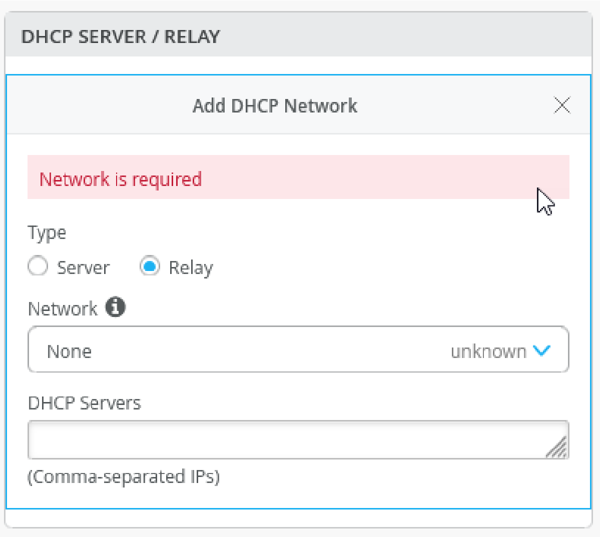

DHCP SERVER / RELAY

Add DHCP Network

Network is required

Type

○ Server ● Relay

Network

None unknown

DHCP Servers

(Comma-separated IPs)

5. From the Network drop-down list, select a network for the DHCP relay.

6. In the DHCP Servers field, configure the IP address for the remote DHCP server. You can add up to three IP addresses in a comma separated format.

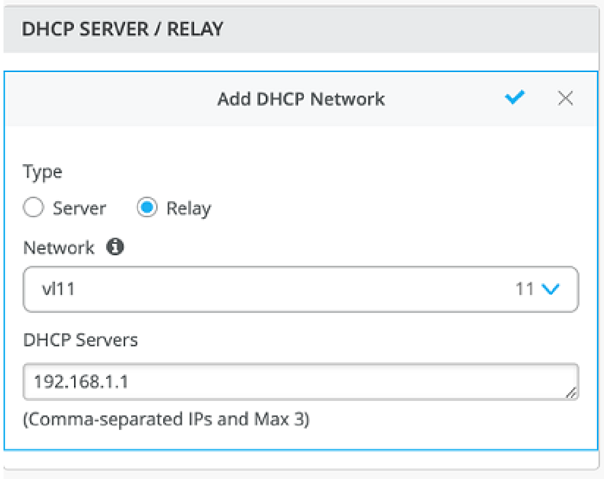

DHCP SERVER / RELAY

Add DHCP Network

Type

○ Server ● Relay

Network

vl11 11

DHCP Servers

192.168.1.1

(Comma-separated IPs and Max 3)

7. Save the changes.

The following configuration will be applied to the API `/sites/:site-id/devices/device-id`:

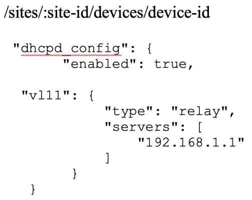

```
/sites/:site-id/devices/device-id

 "dhcpd_config": {
        "enabled": true,

  "vl11": {
             "type": "relay",
             "servers": [
                  "192.168.1.1"
             ]
        }
   }
```

Check the corresponding intended configuration to be pushed on following API: `/sites/:site-id/devices/device-id/config_cmd`

The configuration looks like the below:

```
"set groups top forwarding-options dhcp-relay server-group v111 192.168.1.1",
"set groups top forwarding-options dhcp-relay server-group v111 interface irb.11",
"set groups top forwarding-options dhcp-relay server-group v111 active server group v111",
```

8. Verify the configuration on the switch using the following CLI:

```
{master:0}
root@ex2300-staging> show configuration |compare rollback 1
[edit groups top forwarding-options]
+    dhcp-relay {
+        server-group {
+            vl11 {
+                192.168.1.1;
+            }
+        }
+        group vl11 {
+            active-server-group vl11;
+            interface irb.11;
+        }
+    }
```

9. Make sure the remote DHCP server is reachable from the switch.

```
{master:0}
root@ex2300-staging> ping 192.168.1.1
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=64 time=7.553 ms|
64 bytes from 192.168.1.1: icmp_seq=1 ttl=64 time=16.166 ms
^C
--- 192.168.1.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.553/11.860/16.166/4.306 ms
```

10. Verify DHCP relay binding on the switch.

```
{master:0}
root@ex2300-staging> show dhcp relay binding detail

Client IP Address:  11.0.0.218
     Hardware Address:             d4:20:b0:83:dd:9c
     State:                        BOUND(RELAY_STATE_BOUND)
     Lease Expires:                2022-06-16 18:46:58 IST
     Lease Expires in:             86384 seconds
     Lease Start:                  2022-06-15 18:46:58 IST
     Last Packet Received:         2022-06-15 18:46:58 IST
     Incoming Client Interface:    irb.11:ge-0/0/8.0
     Server Ip Address:            192.168.1.3
     Server Interface:             none
     Bootp Relay Address:          11.0.0.10
     Session Id:                   15
     Relay Id Length:              31
     Relay Id:
/0x00020000/0x00000583/0x01000000/0x00000000
     Relay Id:
/0x66343a62/0x663a6138/0x3a30363a/0x39633a
```

Here are some useful Junos CLI commands:

- `show dhcp relay binding` or `show dhcp relay binding detail`—To view the DHCP relay binding information.
- `show dhcp server statistics`—To view the DHCP relay messages sent/received statistics.
- `clear dhcp relay binding interface x-x/x/x`—To clear the DHCP client binding for a client on a particular interface.
- `clear dhcp relay binding <address>`—To clear the DHCP client binding based on an IP address or MAC address of the wired client.

See also: DHCP Relay Agent.

Configure Bridge Priority on Switches via Mist

SUMMARY

Follow these steps to identify which bridge gets priority.

The bridge priority value along with the device MAC address forms the bridge ID, which determines the bridge to be elected as the root bridge in a Spanning Tree Protocol (STP) topology. The priority value is useful when two bridges have the same path cost to the root bridge.

To configure bridge priority on a switch:

1. Click **Switches** to go to the list of switches.
2. Click a switch to open it.
   The switch details page appears.
3. Scroll down to the STP Bridge Priority tile in the Management section.
4. Select a value from the Bridge Priority drop-down list.

> **NOTE:**
>
> - If you don't set the bridge priority (or select the value ‘none’), the default value (32,768) is configured.
> - The bridge priority can be set only in increments of 4096 (between 0 and 61,440).
> - We recommend that you set the lowest priority for the switch to be elected as root bridge.

5. To save the changes, click **Save** on the upper right of the switch details page.

OSPF Configuration for Switches

SUMMARY

Follow these steps and example to configure OSPF at the organization and site level.

IN THIS SECTION

OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions, making route calculations using the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS or area that contain information about that router's attached interfaces and routing metrics. Each router uses the information in these link-state advertisements to calculate the least cost path to each network and create a routing table for the protocol.

You can also include routing policies (import and export policies) in OSPF configurations at the switch level. The routing policy is composed of terms. Each term can include a set of conditions and a then statement, which defines the actions to take if an OSPF route matches the conditions specified in the term.

Junos OS supports OSPF version 2 (OSPFv2) and OSPF version 3 (OSPFv3), including virtual links, stub areas, and for OSPFv2, authentication. Junos OS does not support type-of-service (ToS) routing.

Example: Configure Basic OSPF in two EX Devices

To configure OSPF in a switch (configure these steps in both the switches):

1. Navigate to the switch by clicking **Switches** > *Switch Name*.
   The switch details page appears.
2. Configure network by following the steps below
   a. Navigate to the NETWORKS tile, and click Add Network.

The New Network window appears.

b. Enter a network name (for example **vlan20**), VLAN ID, and subnet.

c. Click the check mark at the upper right of the New Network window to save the configuration.

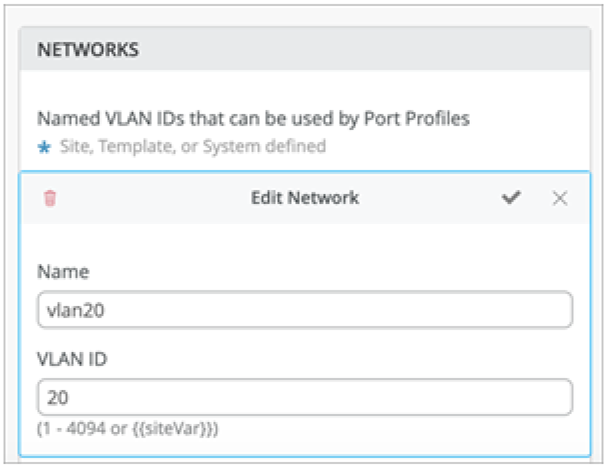

NETWORKS

Named VLAN IDs that can be used by Port Profiles

★ Site, Template, or System defined

Edit Network ✓ ✕

Name

vlan20

VLAN ID

20

(1 - 4094 or {{siteVar}})

3. Create a port profile by following the steps below:

a. On the PORT PROFILES window, click **Add Profile**.
The New Port Profile window appears.

b. Enter the profile details. In this example, name this profile as **vlan20portprofile**. In the port profile, you must include the network (**vlan20**) that you created in the previous step.

c. Click the check mark at the upper right of the New Port Profile window to save the configuration.

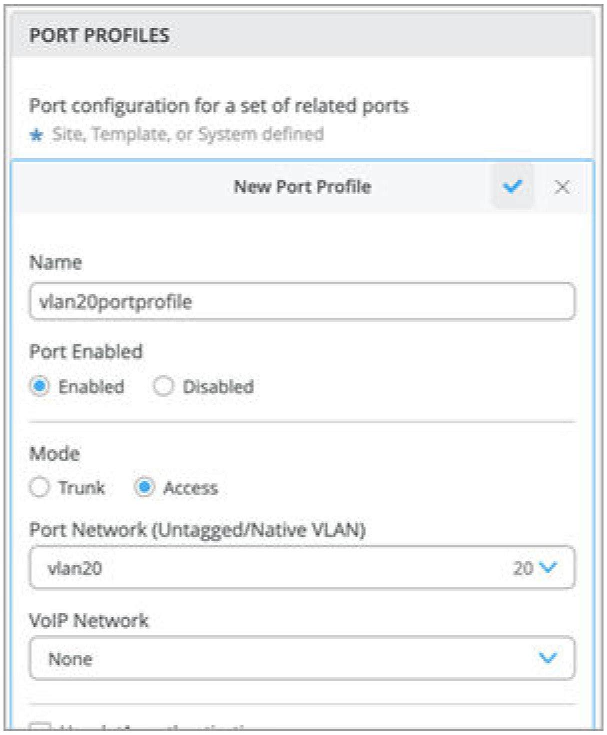

PORT PROFILES

Port configuration for a set of related ports

★ Site, Template, or System defined

New Port Profile ✓ ✕

Name

vlan20portprofile

Port Enabled

● Enabled ○ Disabled

Mode

○ Trunk ● Access

Port Network (Untagged/Native VLAN)

vlan20 20

VoIP Network

None

4. Attach the port profile to the port that is connected to the other switch. To do that:

   a. On the PORT CONFIGURATION tile, click **Add Port Configuration**.
   The New Port Configuration window appears.

   b. Enter the configuration details. Remember to specify the relevant interface in the Port ID field and include the vlan20portprofile in the Configuration Profile field.

   c. Click the check mark at the upper right of the New Port Configuration window to save the configuration.

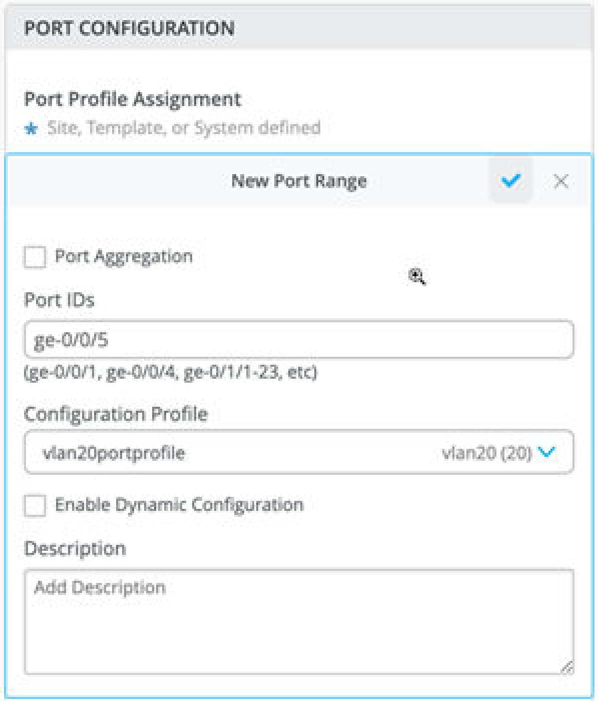

PORT CONFIGURATION

Port Profile Assignment

★ Site, Template, or System defined

New Port Range

□ Port Aggregation

Port IDs

ge-0/0/5

(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Configuration Profile

vlan20portprofile vlan20 (20)

□ Enable Dynamic Configuration

Description

Add Description

5. Add an IP address to the network which you created earlier (vlan20). To do that:

   a. On the IP CONFIGURATION tile, click **Add IP Configuration**.
   The New IP Configuration window appears.

   b. Select Static as address type.

   c. Specify an IP address and subnet mask. In this case, let's use 20.1.1.1/24 on one switch and 20.1.1.2/24 on the other switch.

   d. From the Network (VLAN) drop-down list, select the network (vlan20) that you configured earlier.

   e. Click the check mark at the upper right of the New IP Configuration window to save the configuration.

6. On the OSPF tile, configure an OSPF area by following the steps below:

   a. Click **Add Area**.

   The New Area window appears.

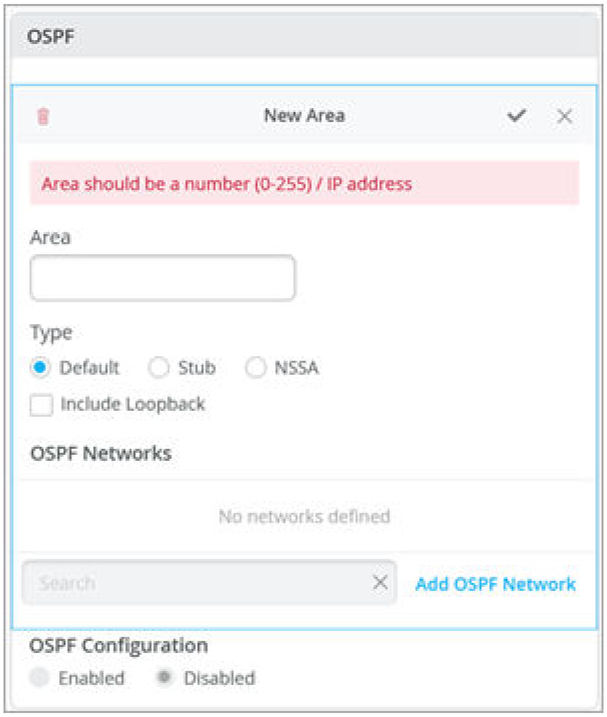

OSPF

New Area

Area should be a number (0-255) / IP address

Area

Type

● Default ○ Stub ○ NSSA

□ Include Loopback

OSPF Networks

No networks defined

OSPF Configuration

○ Enabled ● Disabled

b. Specify the following details:

- Area—Enter the area number (Range: 0 to 255).
- Type—Specify an area type from the drop-down list.
- Include Loopback—Select this check box if you want to include the loopback interface in the OSPF area.

c. Click **Add OSPF Network**.

The Add OSPF Network window appears.

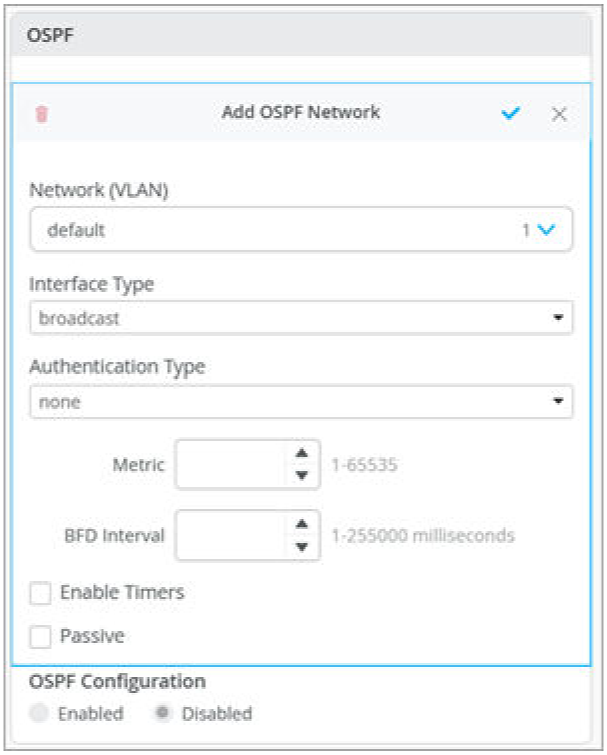

OSPF

Add OSPF Network

Network (VLAN)

default 1

Interface Type

broadcast

Authentication Type

none

Metric 1-65535

BFD Interval 1-255000 milliseconds

□ Enable Timers

□ Passive

OSPF Configuration

○ Enabled ● Disabled

d. Specify the network and the additional details as listed below:

- Network (VLAN)—From the drop-down list, select the network (vlan20) that you defined earlier.

- Interface Type—Select an interface type. The following options are available: broadcast, p2p, and p2mp.

- Authentication Type—Choose an authentication type from the following options: md5,password, and none. If you choose none, ensure that you select that option in both the switches.

- Key—(Applicable if the Authentication Type chosen is md5). Specify a key for md5 authentication. These should be the same in both the devices.

- Value—(Applicable if the Authentication Type chosen is md5). Specify a value for the specified md5 key. These should be the same in both the devices.

- Password—(Applicable if the Authentication Type chosen is password). The password should be the same in both the switches for the OSPF neighborship to be up.

- Metric—Specify the cost of an OSPF interface.

- BFD Interval—Specify the interval at which the device exchanges BFD packets with its peer. Range: 1 through 255000 (in milliseconds).

- Enable Timers—This option allows you to configure Hello Interval and Dead Interval.
- Hello Interval—(Applicable if Enable Times is selected) Specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. By default, the routing device sends hello packets every 10 seconds.
- Dead Interval—(Applicable if Enable Times is selected) Specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. By default, the routing device waits 40 seconds (four times the hello interval).
- Passive—Select this check box to advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.

e. Click the check mark on the upper right of the Add OSPF Network window to save the OSPF Network information.
You are returned to the New Area window.

f. Click the check mark on the upper right of the New Area window to save the OSPF Area information.

7. On the OSPF tile, select the **Enabled** check box under OSPF Configuration.

   When you enable the OSPF configuration, you get additional options to do the following:

   a. Select routing policies (an export policy and an import policy) for the OSPF routes. If a predefined policy does not exist, click **Create Policy** and define a policy by following instructions in "Configure Routing Policies on Switches via Mist" on page 94.

   b. Set a **Reference Bandwidth**, which is used in calculating the default interface cost.

8. On the Routing tile, enter the router ID.
9. Click **Save** on the upper right of the switch details page to save the configuration on the switch.

Carry out the above steps on both the switches to establish the OSPF neighborship between them.

CLI Commands

When Interface Type is p2p and Authentication Type is md5:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 authentication md5 1 key $9$cxXyK8ws...",
"set groups top protocols ospf area 0 interface irb.20 interface type p2p"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2p and Authentication Type is none:

```


"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 interface type p2p"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2p and Authentication Type is password:

```


"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 authentication simple-password
$9$cxXyK8ws...",
"set groups top protocols ospf area 0 interface irb.20 interface type p2p"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2mp and Authentication Type is md5:

```


"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 authentication md5 1 key $9$cxXyK8ws...",
"set groups top protocols ospf area 0 interface irb.20 interface type p2mp"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2mp and Authentication Type is none:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 interface type p2mp"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2mp and Authentication Type is password:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 authentication simple-password
$9$cxXyK8ws...",
"set groups top protocols ospf area 0 interface irb.20 interface type p2mp"
"set groups top routing-options router-id 11.1.1.1",
```

OSPF Events

The following image shows examples of OSPF events:

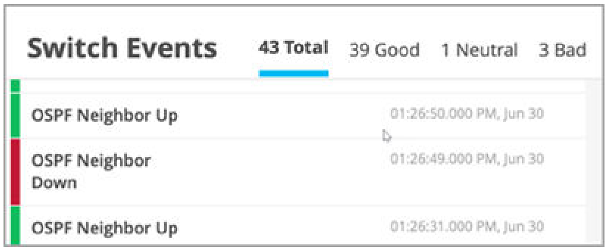

Switch Events

OSPF Neighbor Up 01:26:50.000 PM, Jun 30

OSPF Neighbor Down 01:26:49.000 PM, Jun 30

OSPF Neighbor Up 01:26:31.000 PM, Jun 30

Configure OSPF at the Organization Level

To configure OSPF at the organization template level:

1. Click **Organization** > **Switch Templates**.
2. Open the organization template you want to modify.
3. Navigate to the OSPF tile and click **Add Area**.
4. Follow the steps listed in "Example: Configure Basic OSPF in two EX Devices" on page 137.

Configure OSPF at the Site Level

To configure OSPF at the site template level:

1. Click **Site** > **Switch Configuration**.
2. Open the site template you want to modify.
3. Navigate to the OSPF tile and click **Add Area**.
4. Follow the steps listed in "Example: Configure Basic OSPF in two EX Devices" on page 137.

Add a VRRP Group to a Configuration

SUMMARY

Add one or more VRRP groups to your switch configuration to enable switches to take over for one another in the event of a failover.

IN THIS SECTION

VRRP Overview

Virtual Router Redundancy Protocol (VRRP) is used in high availability (HA) configurations to create a virtual device group, whose member devices share a common virtual IP address. It acts as a virtual switch that acts as the default gateway for the group.

By adding VRRP groups in your configuration, you enable your switches to seamlessly take over for one another in the event of a failover. If a switch in the group is forwarding packets and fails, another switch is selected to automatically replace it.

You can configure VRRP groups directly from the Switch template, site template, or at the device level.

Configure a VRRP Group for a Switch

Before You Begin: In the Shared Elements section of the switch template or site-level switch configuration, add any networks that you want to associate with your VRRP groups. For help with networks, see "Shared Elements" on page 48.

To configure a VRRP group for a switch:

1. Navigate to a switch template, a site-level switch configuration, or a device configuration.

   > **NOTE**: If you need help getting started with a template, site-level switch configuration, or device configuration, see "Overview of Template-Based Switch Configuration" on page 15.

2. Scroll to the **All Switches Configuration** section of the configuration.
3. In the **VRRP** section, click **Add Group**.
4. In the **Add VRRP Network** section:

   a. Enter the group number and authentication type.

   b. Click **Add VRRP Network**, select the Network (VLAN), and enter a VRRP virtual IP address.

   c. Click the check mark to add the VRRP network.

   d. Add more networks if needed.

   e. Click the check mark to add the new VRRP group.

   For help with the various VRRP Group attributes and VRRP Network attributes, see "Attributes for VRRP Groups and VRRP Networks" on page 147.
5. Click **Save** at the top-right corner of the configuration page.

   You can repeat this procedure to add more VRRP groups.

After adding VRRP groups, go to the device configuration for each switch, enable VRRP, and select one or more VRRP groups for the device. You must complete this task for the VRRP settings to be published to the device.

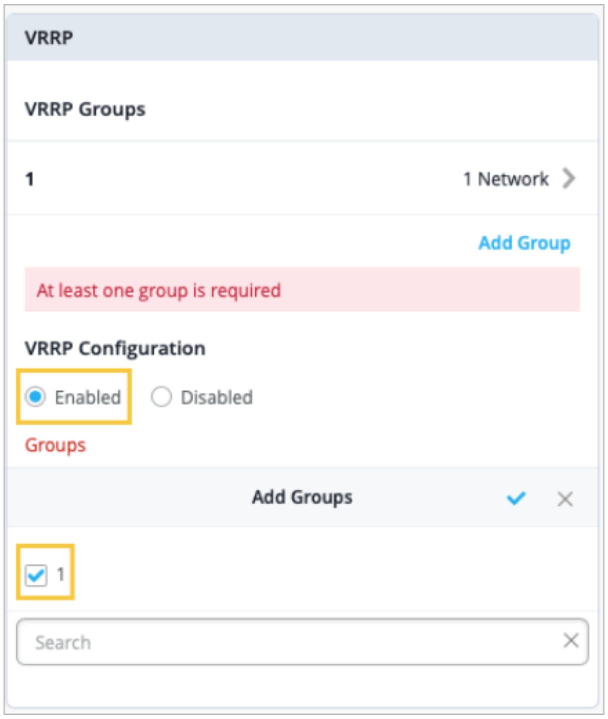

VRRP

VRRP Groups

1 — 1 Network

At least one group is required

VRRP Configuration

● Enabled ○ Disabled

Groups

☑ 1

Attributes for VRRP Groups and VRRP Networks

Table 9: Attributes

Field	Description
Group	Enter a number to identify this group. Range: 0-255
Authentication Type	Select a type from the list. Then enter the settings for the selected type. • md5—If you selected this authentication type, enter these settings: • Key—Enter a number to identify this key. Range: 0-255 • Value—Enter the value for this key. Range: 8-64 • Password—If you selected password, enter a 1- to 8-character password.

Table 9: Attributes *(Continued)*

Field	Description
VRRP Networks	In this section of the New VRRP Group settings, add one or more VRRP networks. Network settings: • **Network (VLAN)**—Select one of the networks that you previously configured in the Networks section of the configuration. For help with networks, see "Shared Elements" on page 48. • **VRRP Virtual IP Address**—You can enter an IP address or use variables. Format: xxx.xxx.xxx.xxx or {{siteVar}}.xxx.xxx

Configure Q-in-Q Tunneling on a Switch Port

SUMMARY

Follow these steps to configure Q-in-Q tunneling to segregate or bundle customer traffic into fewer VLANs or different VLANs.

You can configure switch ports with Q-in-Q tunneling that uses all-in-one bundling. Q-in-Q tunneling enables Layer 2 protocol tunneling (L2PT) on interfaces that are not encapsulation tunnels, and utilizes MAC address rewrite operation.

Using Q-in-Q tunneling, providers can segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q VLAN tags are prepended by the service VLAN (S-VLAN) tag.

You can configure Q-in-Q tunneling under Port Configuration at the switch level, site level, or organization level. The configuration includes selecting Q-in-Q as Configuration Profile, choosing a Port Network (S-VLAN), and specifying other port configuration parameters.

To configure Q-in-Q tunneling in a switch port:

1. Click **Switches** and select.
2. From the **Site** drop-down list, select a site.

   The list of switches in the selected site appears.
3. Click the switch on which you want to configure Q-in-Q tunneling.

   The switch details page appears.
4. On the **Port** tile in the Device section, click **Add Port Configuration**.

   The New Port Configuration window appears.
5. From the Configuration Profile drop-down list, select **Q-in-Q**.
6. Specify the ports on which you want to enable Q-in-Q tunneling, along with other details. The following are the key configuration fields:

   - Port IDs—Specify the port IDs on which you want to configure Q-in-Q tunneling.
   - Interface—Select **L2 interface** as the interface type.
   - Port Network (S-VLAN)—Specify a service VLAN (S-VLAN) if the port is using Q-in-Q tunneling. S-VLAN is an external, additional VLAN tag used to extend Layer 2 Ethernet connections between customer sites. This is especially useful when customers have overlapping VLAN IDs.

   For descriptions of all the fields available on the New Port Configuration window, refer to Table 7 on page 62.
7. Select the check mark on the upper right of the New Port Configuration window to add the port configuration.
8. Click **Save**.

You can configure Q-in-Q tunneling via the organization level templates (Organization > Switch Templates) and site level templates (Organization > Site Configuration) as well. The settings are available on the **Port Config** tab accessed by clicking the **Add Rule** option in the Select Switches Configuration section of the templates.

Manage or Update Configuration Settings

SUMMARY

IN THIS SECTION

Follow these steps to manage templates, update switch settings, and add or delete a CLI configuration.

Manage Templates Settings

The mist portal provides you options to modify, clone, export, or delete a template. If you modify a template, configurations of all the switches managed by that template are modified. You can use the **Export** option to download the template settings in JSON format. You can store the JSON file in your local machine and use it to quickly create new templates, using the **Import** option on the template creation page.

To modify a template:

1. Click **Organization** > **Switch Templates**.
2. Click the template you want to modify. The template opens.
3. Modify the settings. For field descriptions and additional information, refer to "Create a Switch Configuration Template" on page 22.
4. After modifying the template settings, click **Save**.

To clone a template:

1. Click **Organization** > **Switch Templates**.
2. Click the template you want to clone. The template opens.
3. Click **More** > **Clone**.
4. Enter a template name and then click **Clone**. A new template, based on the selected template, is created.

To export a template:

1. Click **Organization** > **Switch Templates**.
2. Click the template you want to export. The template opens.
3. Click **More** > **Export**. The template is downloaded in a JSON file.

To delete a template:

1. Click **Organization** > **Switch Templates**.
2. Click the template you want to delete. The template opens.

3. Click **Delete Template** on the top right.

4. On the Confirm Delete window, click **Delete**. The deleted template is removed from the template list and from all the sites to which it was assigned.

Update Switch Configuration Settings at the Site Level

After applying a template to a site, you can:

- Customize or edit the settings applied to a particular site, if required. You can also replace or unlink a template from the site configuration page.

- Configure additional switch-specific settings from a switch page. The switch-specific settings include a switch name, role, management interface (out of band), and an IRB interface.

To edit the site-level switch configuration settings:

1. Click **Site** > **Switch Configuration**.
2. Click a site from the list to open it.
3. If you want to replace the entire template, select the desired template from the **Configuration Template** drop-down list. If you select the value **(none)**, the existing template gets unlinked from the site.
4. To edit specific template settings of the site:

   a. Select the **Override Configuration Template** in relevant configuration tile.

   b. Edit the settings and then click **Save**. The changes are immediately applied to the switches in the site. For more information, see "Create a Switch Configuration Template" on page 22.

Add or Delete a CLI Configuration

The CLI command options on the switch configuration pages in the Juniper Mist™ portal let you configure features that the predefined drop-down lists and text fields on the Mist portal do not support.

You can add a CLI configuration to a switch by using the `set` command through the Mist portal. Example: `set system ntp server 192.168.3.65`.

Similarly, you can remove a CLI configuration from a switch by using the delete command through the Mist portal. Example: `delete system ntp server 192.168.3.65`.

To add or delete a CLI configuration:

1. Click **Switches** to go to the Switches page.
2. Click your switch from the list to open the switch configuration page.
3. Navigate to the relevant CLI commands box (Additional CLI commands).

   You can also make CLI changes in the **Site/Template CLI Commands** and **Rule-based CLI Commands** boxes available through the switch templates (**Organization** > **Switch Templates**).
4. To add a CLI configuration, enter the set command. For example, `set system ntp server 192.168.3.65`.
5. To remove a configuration, replace set with delete in the command. For example, `delete system ntp server 192.168.3.65`.

The following image shows the delete operation.

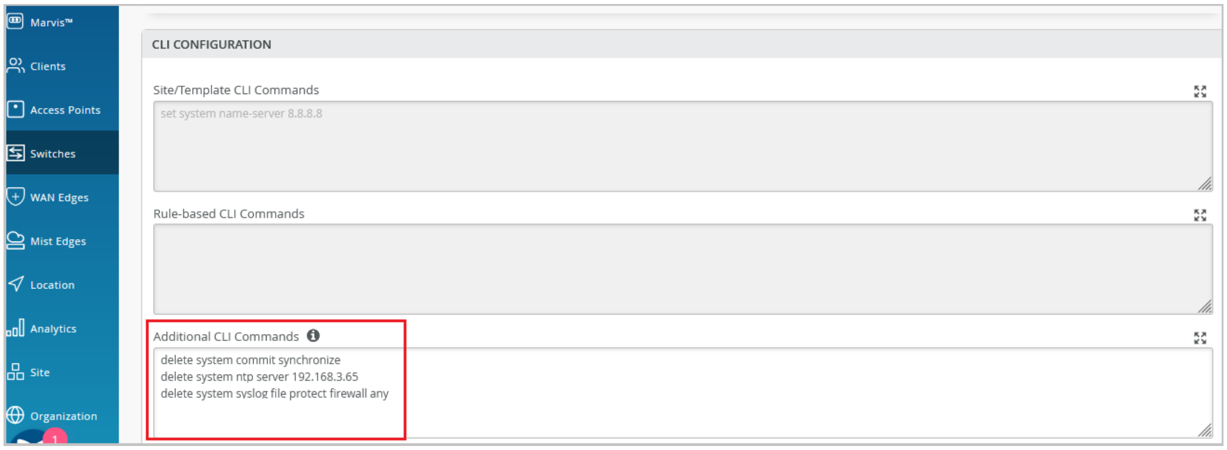

CLI CONFIGURATION

Site/Template CLI Commands

```
set system name-server 8.8.8.8
```

Rule-based CLI Commands

Additional CLI Commands

```
delete system commit synchronize
delete system ntp server 192.168.3.65
delete system syslog file protect firewall any
```

When you save the `delete` commands, the following operations take place:

- Mist sends the `delete` commands to the switch.
- The Switch Insights page on the Mist portal generates a **Config Changed by User** event, with a response **UI_COMMIT_COMPLETED**. You can access Switch Insights from "Switch Details" on page 294.
- Mist deletes the CLI commands from the switch. Later, if required, you can remove these commands from the CLI commands box on the Mist portal.
- Mist updates the `delete` commands in the following API call:

```
https://api.mist.com/api/v1/sites/<site_id>/devices/00000000-0000-0000-1000-<switch_mac>/
config_cmd
```

Just selecting a few commands from any CLI command box on the Mist portal and hitting the Backspace or Delete button does not remove the commands from the switch. It removes the commands only from the API, which contains the current switch configuration that is present on the Mist portal.

Deleting the CLI commands only from the GUI also generates a **Config Changed by User** event on the Switch Insights page. However, this event doesn't show the **UI_COMMIT_COMPLETED** response. The changes are made only on the Mist portal GUI, not on the switch.

We don't recommend logging in to the switch CLI and making any changes there if the Mist cloud manages your switch. The changes you make on a switch through the CLI don't get included in the switch configuration in the Mist cloud.

Upgrade Junos OS on Switches

IN THIS SECTION

You can either manually upgrade switches and Virtual Chassis in a site, or schedule or automate upgrades across multiple sites in your organization.

NOTE:
Wired Assurance is not compatible with Junos Flex images. To ensure that only a standard Junos image is deployed, we recommend that you upgrade switches through the Mist portal.

Upgrade Switches in a Site Manually

IN THIS SECTION

NOTE: To schedule or automate switch upgrades across multiple sites in your organization, refer to the instructions provided in "Schedule and Automate Switch Upgrades" on page 162.

Users with Super User or Network Admin privileges for the site can manually "upgrade Junos " on page 157 on switches or Virtual Chassis by selecting them from the switch list and clicking the **Upgrade Firmware** button.

Figure 4: Upgrade Juons on Selected Switches

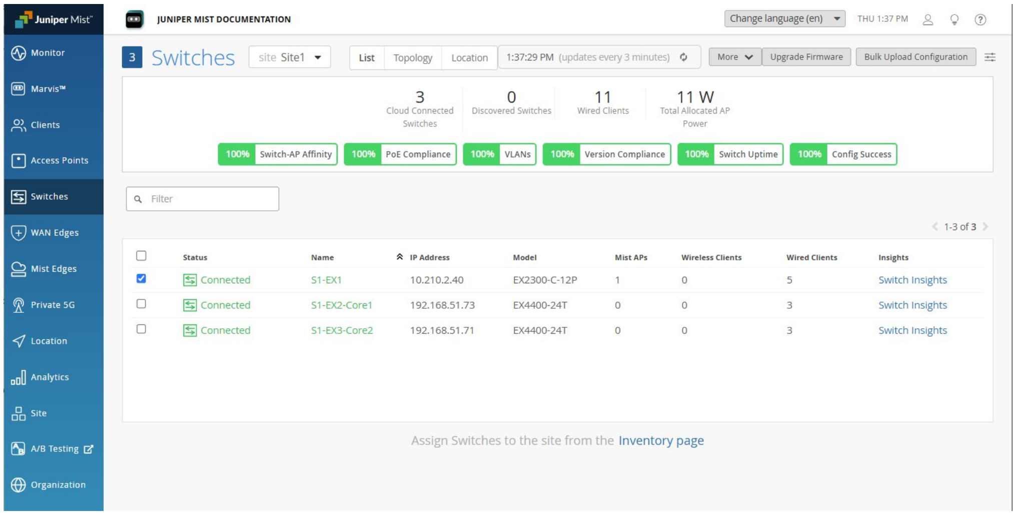

3	0	11	11 W
Cloud Connected Switches	Discovered Switches	Wired Clients	Total Allocated AP Power

100% Switch-AP Affinity	100% PoE Compliance	100% VLANs	100% Version Compliance	100% Switch Uptime	100% Config Success

□	Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	Insights
☑	Connected	S1-EX1	10.210.2.40	EX2300-C-12P	1	0	5	Switch Insights
□	Connected	S1-EX2-Core1	192.168.51.73	EX4400-24T	0	0	3	Switch Insights
□	Connected	S1-EX3-Core2	192.168.51.71	EX4400-24T	0	0	3	Switch Insights

The switch should be under warranty, have an active maintenance contract, and an active software subscription. In addition, please be sure the switch has the following:

- The storage space required to accommodate the new image.
- A stable SSH connection to the Mist cloud.
- (Optional) A recovery snapshot stored on the OAM volume. See "Switch Utilities" on page 307 for details about the snapshot.

Free Up Storage Space on Your Switch

When you initiate a switch upgrade process, Juniper Mist™ runs the `request system storage cleanup` command on the switch before copying the software image. This process mostly ensures the availability of storage space to accommodate the software image in the /var/tmp folder on the switch. However, in the case of some switches, **such as EX2300 and EX3400**, the `request system storage cleanup` command doesn't clear the required space. In this case, you will need to free up more space.

NOTE:

- To perform the steps listed in this topic, you must have the root password configured in the site settings on the **Organization** > **Site Configuration** page of the Juniper Mist portal.
- Perform the steps listed in this topic only if your switch doesn't have the required space for the upgrade.

To free up storage space on your switch:

1. On the Juniper Mist portal, click **Switches** to go to the list of switches.
2. Locate the switch on which you want to perform the storage cleanup operation.
3. Select **Utilities** > **Remote Shell**.
4. Begin a shell session by entering the `start shell user root` command, followed by the root password.

   ```
   {master: 0}
   mist@Mist_Sw> start shell user root
   Password:
   root@Mist_Sw:RE:0%
   ```

   This step starts a shell session on the primary FPC member by default.
5. Check the storage usage, by running the `df -h` command.

   Generally, the `/dev/gpt/junos` file system takes up most of the space.

   ```
   user@Mist_Sw:RE:0% df -h
   Filesystem          Size    Used      Avail     Capacity     Mounted on
   /dev/md0.uzip        22M     22M         0B          100%        /
   devfs               1.0K    1.0K         0B          100%        /dev
   /dev/gpt/junos      1.3G    941M       315M           75%        /.mount
   ...output truncated...
   ```
6. Run the following command to free up the space on the switch:

   ```
   root@Mist_Sw :RE:0% pkg setop rm previous
   root@Mist_Sw :RE:0% pkg delete old
   ```

7. Check the available storage, using the `df -h` command. The output now shows lesser space as used under `/dev/gpt/junos`.

```
user@Mist_Sw:RE:0% df -h
Filesystem          Size    Used      Avail     Capacity     Mounted on
/dev/md0.uzip        22M     22M         0B         100%        /
devfs               1.0K    1.0K         0B         100%        /dev
/dev/gpt/junos      1.3G    567M       689M          45%        /.mount
...output truncated...
```

8. Exit the shell session to return to the CLI operational mode, and then check the storage usage from there.

```
user@Mist_Sw:RE:0% exit
exit

{master :0}

user@Mist_Sw> show system storage

Filesystem          Size    Used      Avail     Capacity     Mounted on
/dev/gpt/junos      1.3G    941M       315M          75%        /.mount
tempfs              393M     68K       393M           0%        /.mount/tmp
tempfs              324M    576K       324M           0%        /.mount/mfs
...output truncated...
```

In the case of a Virtual Chassis upgrade, the preceding steps free up the space only on the primary member (member 0). You also need to initiate a session with each of the other FPC members (such as member 1 and member 2) and repeat the storage cleanup steps. See the following example:

```
user@Mist_Sw> request session member 1
Last login: Tue Feb 16 00:42:30 from 13.56.90.212...

mist@Mist_Sw> start shell user root
Password:

user@mist_sw:RE:0% df -h
Filesystem          Size    Used      Avail     Capacity     Mounted on
```

```
/dev/md0.uzip        22M     22M         0B         100%        /
devfs               1.0K    1.0K         0B         100%        /dev
/dev/gpt/junos      1.3G    916M       340M          73%        /.mount
...output truncated...
```

Upgrade the Switch

Supported Devices

The Juniper Mist™ portal supports upgrading the Junos OS software on the following platforms: EX2300, EX3400, EX4000, EX4100, EX4100-F, EX4300-P, EX4300-MP, EX4400, EX4600, EX4650, EX9200, QFX5110, QFX5120, and EX Series Virtual Chassis.

Juniper Mist does not support nonstop software upgrade (NSSU).

Available Versions

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For example, you can upgrade from 21.2 to the next three releases—21.3, 21.4 and 22.1—or downgrade to the previous three releases—21.1, 20.4 and 20.3.

For EEOL releases, you have an additional option—you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 21.2 is an EEOL release. Hence, you can upgrade from 21.2 to the next two EEOL releases —21.4 and 22.2—or downgrade to the previous two EEOL releases—20.4 and 20.2. Check Junos OS Dates and Milestones to see whether a release has reached EEOL.

Selecting a Release

For more information about releases, consult these topics:

- Suggested Releases to Consider and Evaluate
- Knowledge Base (Log in, and then search by service release number.)
- Junos OS Installation and Upgrade Overview
- Junos OS Evolved Installation Packages

Initiating a Switch Upgrade

To upgrade the Junos OS software on your switch (or Virtual Chassis):

1. Click **Switches** on the left navigation pane in the Juniper Mist portal.

2. From the **Site** menu, select the site where you want to perform the upgrade.
3. Locate the switch to be upgraded, and ensure that it is connected (displays the Connected status).

   If the switch doesn't appear on Mist as connected, troubleshoot the issue as explained in "Troubleshoot Your Switch Connectivity" on page 325.
4. From the **List** tab, select the switch that requires a software upgrade, and then click **Upgrade Firmware**. You can select one or more switches for upgrade.

   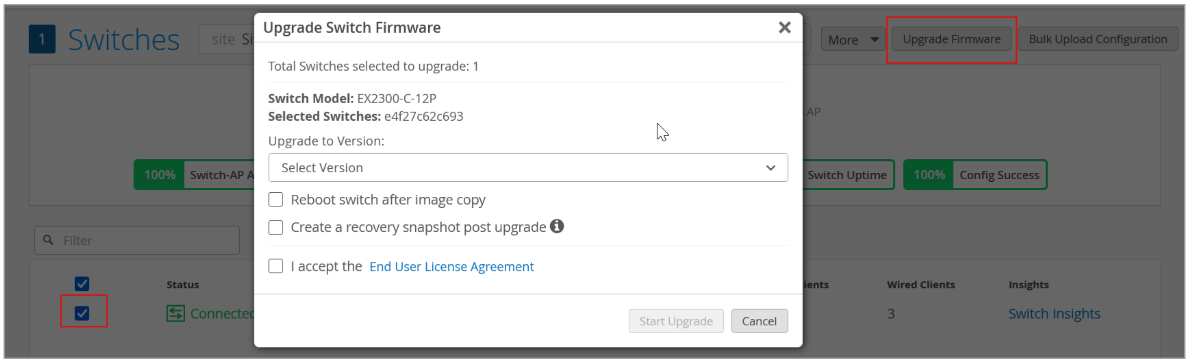

   Upgrade Switch Firmware ×

   Total Switches selected to upgrade: 1

   Switch Model: EX2300-C-12P
   Selected Switches: e4f27c62c693

   Upgrade to Version:
   Select Version

   □ Reboot switch after image copy
   □ Create a recovery snapshot post upgrade
   □ I accept the End User License Agreement

   Start Upgrade   Cancel

   Alternatively, you can also upgrade the switch by using the **Upgrade Firmware** option on the **Utilities** drop-down list on the switch details page (see "Switch Details" on page 294).
5. In the Upgrade Switch Firmware window, select the target software version from the **Upgrade to Version** drop-down list, and then click **Start Upgrade**. The drop-down list displays the suggested software version for the selected switch, along with all the applicable versions.

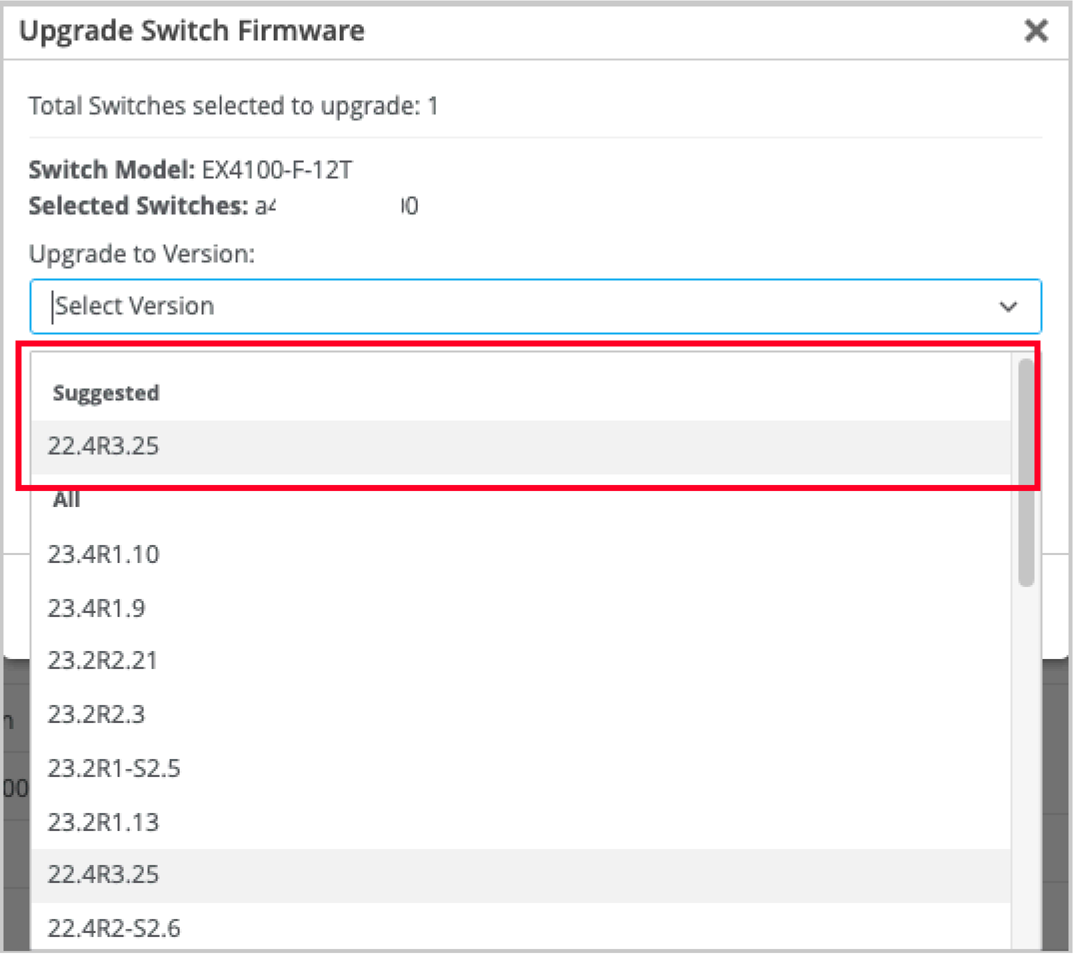

Upgrade Switch Firmware ✕

Total Switches selected to upgrade: 1

Switch Model: EX4100-F-12T
Selected Switches: a4 0

Upgrade to Version:

Select Version

If you don't see the software version you are looking for, write to support@mist.com. We will make the version available from 24 to 48 hours after receiving the request.

Select the **Reboot switch after image copy** check box if you want the switch to reboot automatically after the image copy procedure is complete.

- If you select this option, the switch boots up with the new image.
- If you do not select this option, the switch remains in a state of pending reboot. In this case, do the following to complete the upgrade:

  a. Click **Switches** > ***Switch Name*** to navigate to the switch details page.

  b. Reboot the switch by clicking the Reboot option displayed at the top of the page. Or, click **Utilities** > **Reboot Switch** to initiate a reboot.

     Depending on the switch model, the switch details page may provide options to either reboot the switch or revert the upgrade. If you choose not to proceed with the upgrade, you can revert it. Note that the revert option is only available until the switch is rebooted.

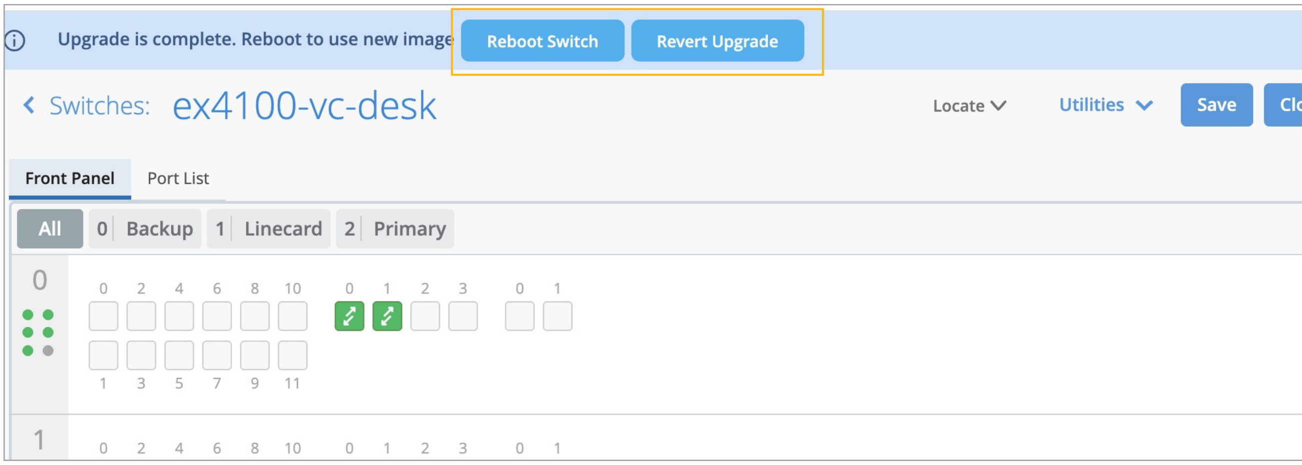

Upgrade is complete. Reboot to use new image

Select **Create a recovery snapshot post upgrade** if you want the switch to have a recovery snapshot. A recovery snapshot stored in OAM (Operations, Administration, and Maintenance) volume holds a full backup that can be used in case something goes wrong with Junos volume.

Select **I accept End User Agreement.**

Once the upgrade starts, the Status column in the switch list view shows the switch status as Upgrading. The column also shows the progress of the upgrade.

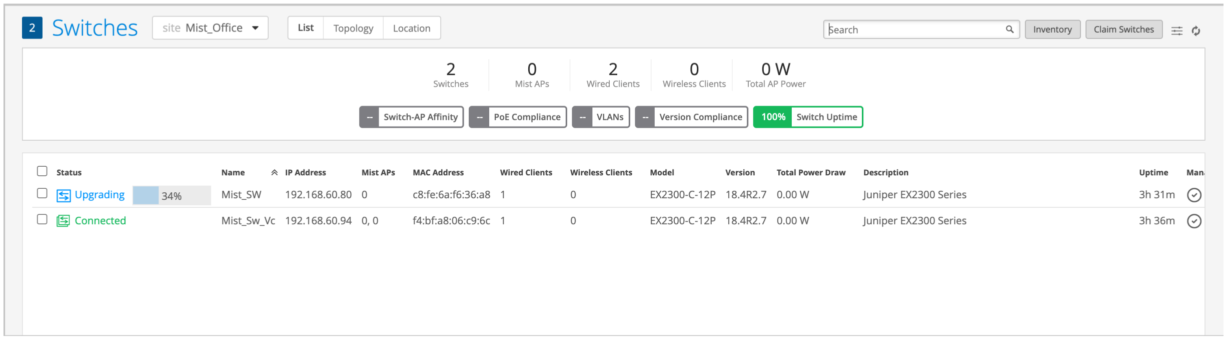

2 Switches | 0 Mist APs | 2 Wired Clients | 0 Wireless Clients | 0 W Total AP Power

-- Switch-AP Affinity | -- PoE Compliance | -- VLANs | -- Version Compliance | 100% Switch Uptime

Status		Name	IP Address	Mist APs	MAC Address	Wired Clients	Wireless Clients	Model	Version	Total Power Draw	Description	Uptime	Man
Upgrading	34%	Mist_SW	192.168.60.80	0	c8:fe:6a:f6:36:a8	1	0	EX2300-C-12P	18.4R2.7	0.00 W	Juniper EX2300 Series	3h 31m	✓
Connected		Mist_Sw_Vc	192.168.60.94	0, 0	f4:bf:a8:06:c9:6c	1	0	EX2300-C-12P	18.4R2.7	0.00 W	Juniper EX2300 Series	3h 36m	✓

If you don't see the Status column in the switch list view, click the hamburger menu in the upper right of the page. Select the Status check box to display the column.

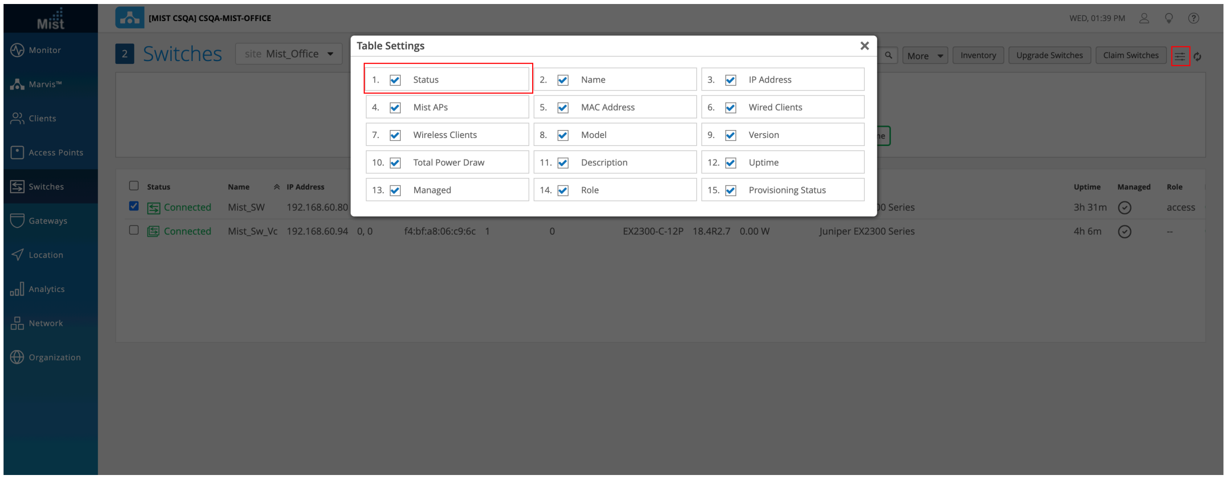

Table Settings

1. ☑ Status
2. ☑ Name
3. ☑ IP Address
4. ☑ Mist APs
5. ☑ MAC Address
6. ☑ Wired Clients
7. ☑ Wireless Clients
8. ☑ Model
9. ☑ Version
10. ☑ Total Power Draw
11. ☑ Description
12. ☑ Uptime
13. ☑ Managed
14. ☑ Role
15. ☑ Provisioning Status

	Status	Name	IP Address									Uptime	Managed	Role
☑	Connected	Mist_SW	192.168.60.80								0 Series	3h 31m	✓	access
□	Connected	Mist_Sw_Vc	192.168.60.94	0, 0	f4:bf:a8:06:c9:6c	1	0	EX2300-C-12P	18.4R2.7	0.00 W	Juniper EX2300 Series	4h 6m	✓	--

You can also view the switch status (as Upgrading) on the switch details page and the Switch Insights page.

You can view the upgrade events in the Switch Events section of a Switch Insights page. To access the Switch Insights page, open a "switch details" on page 294 page and click the **Switch Insights** link on the **Properties** tile.

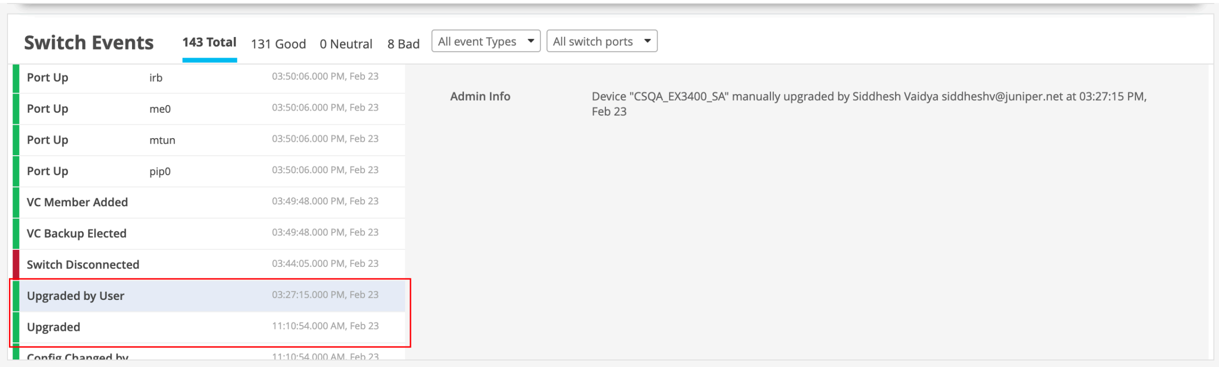

Switch Events 143 Total 131 Good 0 Neutral 8 Bad

Port Up	irb	03:50:06.000 PM, Feb 23
Port Up	me0	03:50:06.000 PM, Feb 23
Port Up	mtun	03:50:06.000 PM, Feb 23
Port Up	pip0	03:50:06.000 PM, Feb 23
VC Member Added		03:49:48.000 PM, Feb 23
VC Backup Elected		03:49:48.000 PM, Feb 23
Switch Disconnected		03:44:05.000 PM, Feb 23
Upgraded by User		03:27:15.000 PM, Feb 23
Upgraded		11:10:54.000 AM, Feb 23
Config Changed by		11:10:54.000 AM, Feb 23

Admin Info Device "CSQA_EX3400_SA" manually upgraded by Siddhesh Vaidya siddheshv@juniper.net at 03:27:15 PM, Feb 23

The above image shows a Switch Insights page, which lists switch upgrade events. The Upgraded by User event indicates that a user has initiated the upgrade. The Upgraded event indicates that the upgrade operation is complete. This means that the new software image was copied and the switch was rebooted.

An upgrade will fail if:

- The switch doesn't have an SSH connection to the Juniper Mist cloud or if an uplink port is flapping.
- The switch doesn't have enough storage. If the upgrade fails because of insufficient space, the upgrade failure event is displayed on the Switch Insights page as shown below:

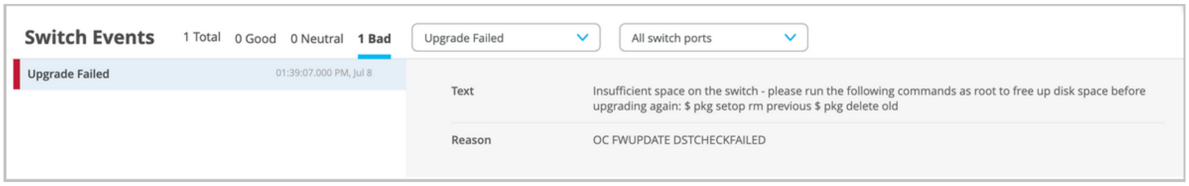

Upgrade Failed 01:39:07.000 PM, Jul 8

Text: Insufficient space on the switch - please run the following commands as root to free up disk space before upgrading again: $ pkg setop rm previous $ pkg delete old

Reason: OC FWUPDATE DSTCHECKFAILED

See also: Free Up Storage Space on Your Switch

- You initiate an upgrade to the same software version that is already running on the switch. In this case, the Switch Events section of the Switch Insights page shows this failure reason:

  Upgrade not needed. Please check current or pending version.

- The time on the switch is incorrect. In this case, the Switch Events section of the Switch Insights page shows this failure reason:

  OC FWUPDATE WRITEFAILED. See also: [EX/QFX] Certificate errors - Cannot validate Junos Image : Format error in certificate.

Schedule and Automate Switch Upgrades

IN THIS SECTION

NOTE: To manually upgrade individual switches or a bunch of switches in a site, refer to the instructions provided in "Upgrade Switches in a Site Manually" on page 153.

Juniper Mist allows you to:

- Create and manage upgrade schedules for the switches that are connected to the Mist cloud. This option is available at the organization and site level.

- Configure settings to automatically upgrade new switches when they are onboarded. This setting is available only at the organization level.

Schedule Switch Upgrades

You can schedule firmware upgrades (Junos OS upgrades) on your switches for a future date and time. You can also execute the upgrades immediately. You can create upgrade schedules at the organization and site level for switch models that are already connected to cloud.

To schedule a switch upgrade:

1. To schedule an upgrade at the organization level:

   a. Click **Organization** > **Settings.**

   b. Navigate to the Firmware Upgrade tile.

   Or, to schedule an upgrade at the site level:

   a. Click **Organization** > **Site Configuration**.

   b. Click a site configuration record to open it.

   c. Navigate to the **Firmware Upgrade** tile and then click the **Switches** tab.

2. Click **Add Upgrade**.

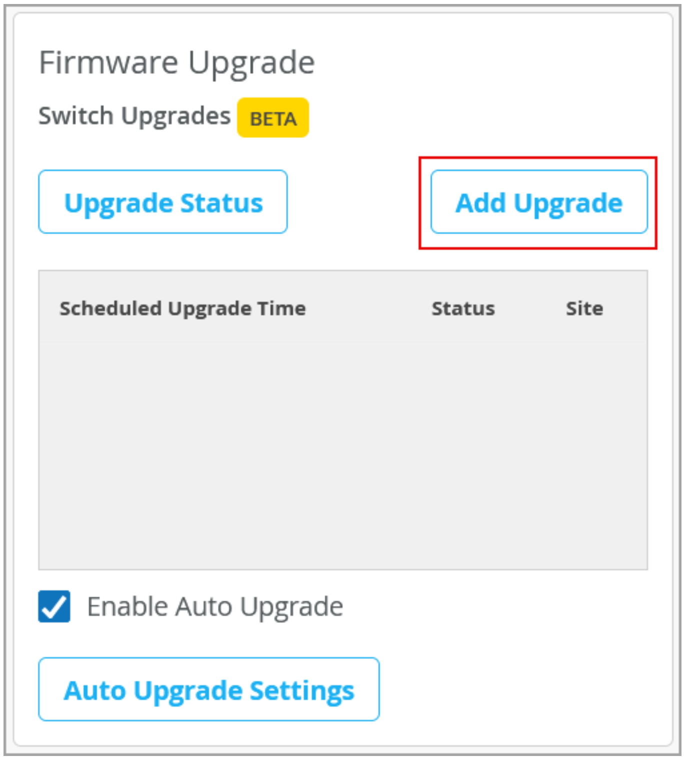

☑ Enable Auto Upgrade

3. (Applicable to site-level upgrade task) Click the **+** button, select the sites for which you need to run the upgrade, and then click **Add**.

4. Click **Next**.

   All the device models available in the inventory are listed.

5. For each device model listed, select a target Junos version from the **Upgrade to Version** drop-down list.

6. If you want to apply the upgrade only to certain specific switches, do the following:

   a. Enable the **Match Criteria** option. This option allows you to define certain additional parameters that need to be matched for a switch to be included in an upgrade schedule.

   b. Define the parameters to be matched. They include:

      - **Switch Name**—Specify a switch name to be matched. Optionally, specify an offset value, which indicates the starting character within the switch name (0 being the first character). For example, to match the keyword **IDF-2** in the switch hostname **JNPR-IDF-2**, enter **JNPR-IDF-2** as the switch name and set the offset value to 5.

      - **Role**—You can either select a switch role that is already defined by a user, or define a new role.

      - **Campus Fabric Role**—Select a campus fabric role from the drop-down list.

> **NOTE**: If multiple parameters (for example, a Switch Name and a Role) are specified, the upgrade is triggered only if all the specified parameters are matched. The keywords are case-sensitive.

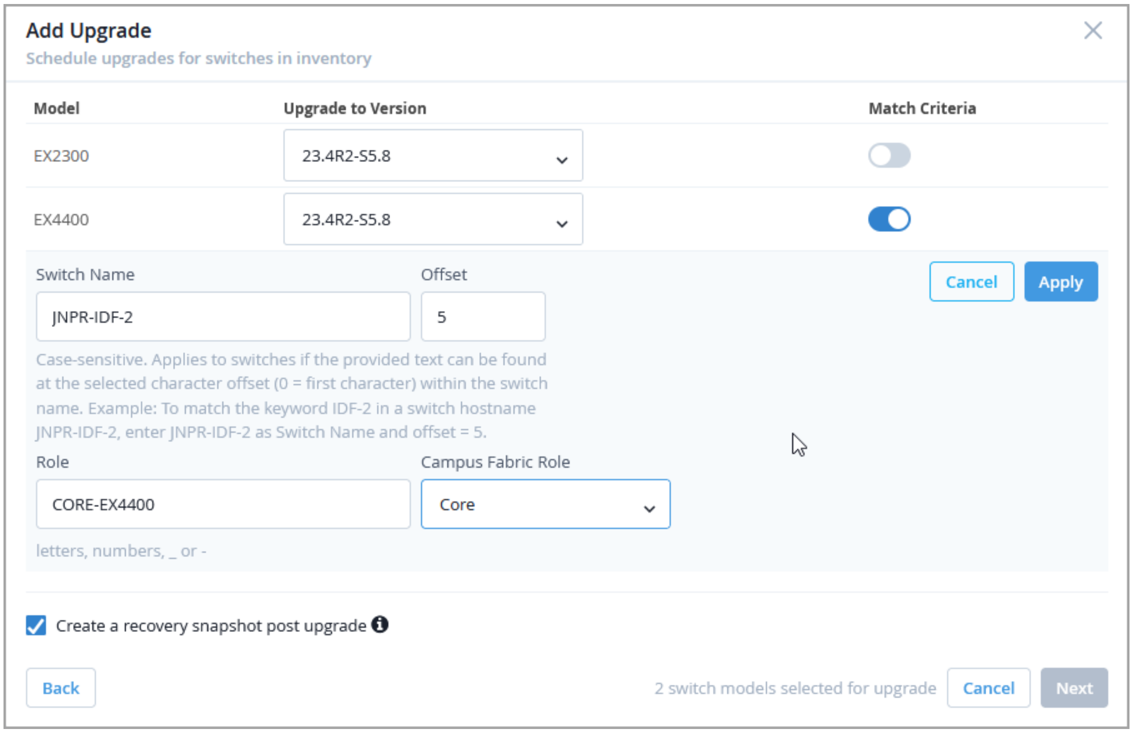

Add Upgrade

Schedule upgrades for switches in inventory

Model	Upgrade to Version	Match Criteria
EX2300	23.4R2-S5.8	○
EX4400	23.4R2-S5.8	●

Switch Name: JNPR-IDF-2

Offset: 5

Case-sensitive. Applies to switches if the provided text can be found at the selected character offset (0 = first character) within the switch name. Example: To match the keyword IDF-2 in a switch hostname JNPR-IDF-2, enter JNPR-IDF-2 as Switch Name and offset = 5.

Role: CORE-EX4400

letters, numbers, _ or -

Campus Fabric Role: Core

☑ Create a recovery snapshot post upgrade

2 switch models selected for upgrade

   c. Click **Apply**.

Match criteria configuration is available per device model.

NOTE: A recovery snapshot is automatically created on the switch post upgrade. If you don't want that, you must clear the **Create a recovery snapshot post upgrade** check box.

7. Click **Next** to go to the tab where you can specify the schedule details.
8. Specify the following information:

Field	Description
Image Installation Schedule	Select one of the following options: • **Image Installation Now**: Choose this option if you want to execute the upgrade immediately. This is the default option. • **Image Installation Later**: Choose this option to schedule the switch upgrade for a future date and time. You need to also specify a date and time using the date picker in the **Image Installation Time** field. You can only specify a date that falls within the next two months.
Image Installation Strategy	Select an image installation strategy. These strategies will be applied separately to each upgrade run (generated for each site-model-SKU combination). The following strategies are available: • Serial: Downloads and installs images on switches in a sequential, random order. • Phased: Downloads and installs images on switches in user-defined phases (specified in percentage). If you select this option, you need to also specify the phases in percentages. For example, to upgrade 100 devices in four phases, you can specify the values 5, 25, 50, 100 in the **Image Installation Canary Phases** field. • Simultaneous: Installs images on all switches at once without any specific priority. This is the default option.
Image Installation Max Failure percentage	Maximum percentage of devices allowed to fail upgrade in each upgrade run (generated for each site-model-SKU combination). If the number of failures exceeds the set percentage: • The upgrade process is cancelled for the remaining devices in the current run. • Reboot is not triggered for any of the devices included in the current upgrade run.

Reboot Devices	Select one of the following options: - Reboot Now: Reboots the device immediately after the image installation process is completed on all the devices scheduled for upgrade. - Reboot Later: Select this option to schedule reboot to a future time. You need to specify a date and time using the date picker in the **Reboot Time** field. You can only specify a date that is within the next two months, but not earlier than the image installation date and time.
Reboot Strategy	Select a reboot strategy. These strategies will be applied separately to each upgrade run (generated for each site-model-SKU combination). The following strategies are available: - Serial: Reboots the switches in a sequential, random order. - Phased: Reboots the switches in user-defined phases. If you select this option, you need to specify the phases in percentages. For example, to reboot 100 devices in four phases, you can specify the values 5, 25, 50, 100 in the **Reboot Canary Phases** field. - Simultaneous: Reboots all devices at once without any specific priority.
Reboot Max Failure percentage	Maximum percentage of devices allowed to fail reboot in each upgrade run (generated for each site-model-SKU combination). If the number of failures exceeds the set percentage, the reboot process is cancelled for the remaining devices in the current run.

9. Select the **I accept the End User License Agreement** check box.
10. Click **Save and Start Upgrade**.

If you configure multiple upgrade schedules for the same device, the most recently configured schedule will be applied.

View Switch Upgrade Status

You can view the scheduled switch upgrades and their statuses on the Firmware Upgrade tile on the organization settings (for organization-level upgrades) and site configuration (for site-level upgrades) pages.

Note that the upgrade statuses will be available only if your inventory has switches that meet the upgrade schedule parameters defined. For example, if you specify a switch name in the upgrade schedule but your inventory does not have a switch with that name, you will not see a status record for that upgrade schedule.

To view the details of all the future and past switch upgrades:

1. To view the details of the upgrades configured at the organization level:

  a. Click **Organization** > **Settings.**

     If the organization has any upgrades scheduled for any sites, you can see a message indicating the same on the upper side of the organization settings page.

  b. Navigate to the **Firmware Upgrade** tile.

     All the scheduled upgrades for the organization along with their statuses are listed on this tile.

  Or, to view the details of the upgrades configured at the site level:

  a. Click **Organization** > **Site Configuration**.

  b. Click a site configuration record to open it.

  c. Navigate to the Firmware Upgrade tile and then click the **Switches** tab. All the scheduled upgrades for the selected site along with their statuses are listed on this tile.

2. Click **Upgrade Status**.

   The **Scheduled Upgrades** tab is displayed.

3. Expand the upgrade schedule that you want to view.

4. On the **Scheduled Upgrades** tab, you can do the following:

   - View the upgrade details.
   - Click the hyperlink in the **Model** column to view the switches included in the selected upgrade schedule.
   - Click **Edit** or **Cancel Upgrade** to modify or cancel the upgrade schedules.
   - Click **Add Upgrade** to create a new upgrade schedule.

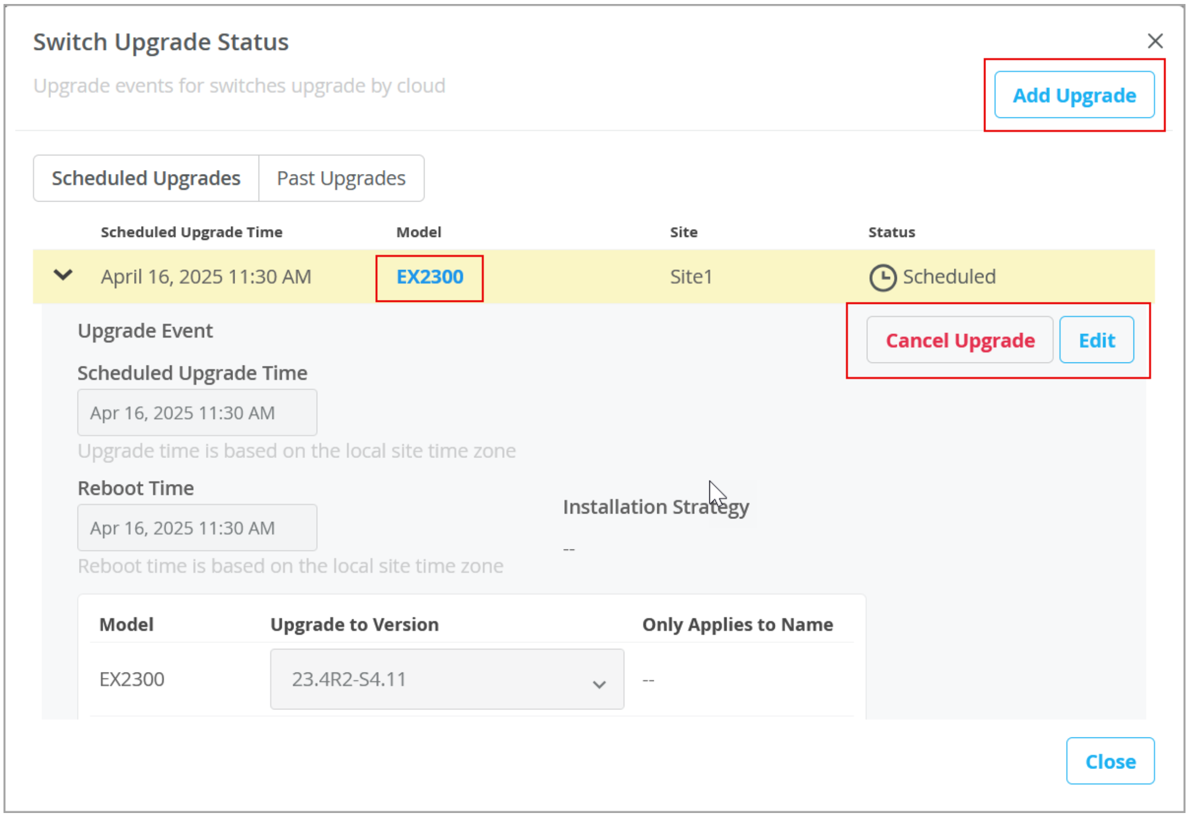

Scheduled Upgrade Time	Model	Site	Status
April 16, 2025 11:30 AM	EX2300	Site1	Scheduled

Upgrade Event

Scheduled Upgrade Time: Apr 16, 2025 11:30 AM

Upgrade time is based on the local site time zone

Reboot Time: Apr 16, 2025 11:30 AM

Reboot time is based on the local site time zone

Installation Strategy: --

Model	Upgrade to Version	Only Applies to Name
EX2300	23.4R2-S4.11	--

To view all the upgrades that were completed within the last 30 days, click the **Past Upgrades** tab. You cannot edit these records.

Modify or Cancel Upgrade Schedules

To modify or cancel any switch upgrade schedules:

1. To modify or cancel an upgrade schedule at the organization level:

   a. Click **Organization** > **Settings.**

   b. Navigate to the Firmware Upgrade tile.

   Or, to modify or cancel an upgrade schedule at the site level:

   a. Click **Organization** > **Site Configuration**.

   b. Click a site configuration record to open it.

   c. Navigate to the Firmware Upgrade tile and then click the **Switches** tab.
2. Click **Upgrade Status**.

   The **Scheduled Upgrades** tab is displayed.
3. Expand the upgrade schedule that you want to update or cancel, and use the following options as required:

- To edit the schedule, click **Edit**, make the changes, and then click **Save Changes**. You can edit only those upgrades that have not yet started. You cannot edit in-progress upgrades. Only the following information can be edited: scheduled upgrade time, reboot time, and the upgrade to version.
- To cancel the upgrade schedule, click **Cancel Upgrade**. You can cancel the upgrades that are either in-progress or yet to start.

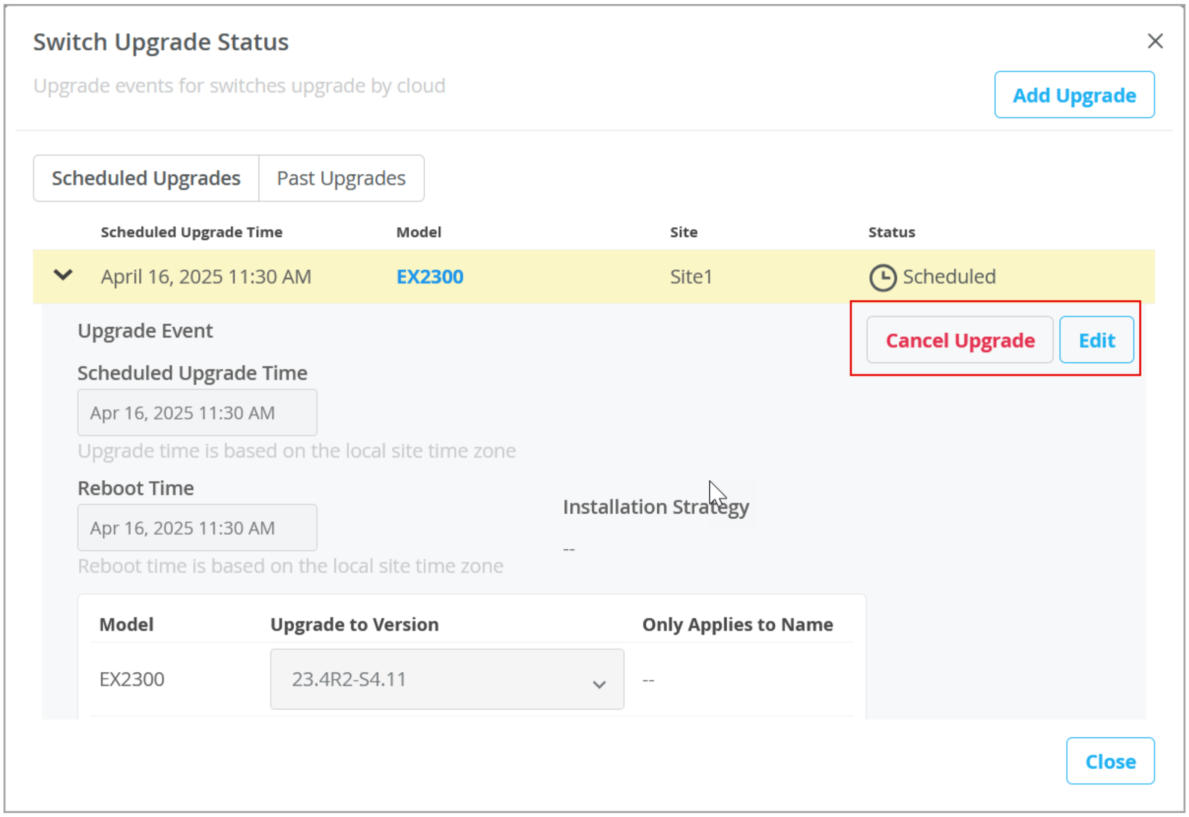
Switch Upgrade Status

Upgrade events for switches upgrade by cloud

Scheduled Upgrade Time	Model	Site	Status
April 16, 2025 11:30 AM	EX2300	Site1	Scheduled

Upgrade Event

Scheduled Upgrade Time: Apr 16, 2025 11:30 AM

Upgrade time is based on the local site time zone

Reboot Time: Apr 16, 2025 11:30 AM

Reboot time is based on the local site time zone

Installation Strategy: --

Model	Upgrade to Version	Only Applies to Name
EX2300	23.4R2-S4.11	--

Enable Automatic Upgrade for Switches

If you want newly onboarded switches in an organization to automatically upgrade to a specific Junos version when they connect to the Mist cloud for the first time, you can map switch models to relevant Junos upgrade versions. The automatic upgrade settings can be configured only at the organization level. This feature is not applicable to the switches that are already online.

NOTE: We recommend enabling Auto Upgrade as a best practice to ensure that all the onboarded switches are on the suggested version right from the beginning.

To configure automatic upgrade settings:

1. Click **Organization** > **Settings.**
2. Navigate to the **Firmware Upgrade** tile.

3. Select the **Enable Auto Upgrade** check box, and then click **Auto Upgrade Settings**.

   The Switch Auto Upgrade Page appears.

4. Select the **I accept the End User License Agreement** check box.

   > **NOTE**: A recovery snapshot is automatically created on the switch post upgrade. If you don’t want that, you must clear the **Create a recovery snapshot post upgrade** check box.

5. Click **Save**.

When you claim a new switch, the claim window displays the automatic upgrade version of Junos that is configured for each switch model. If you want to update this configuration, click **Organization Settings**.

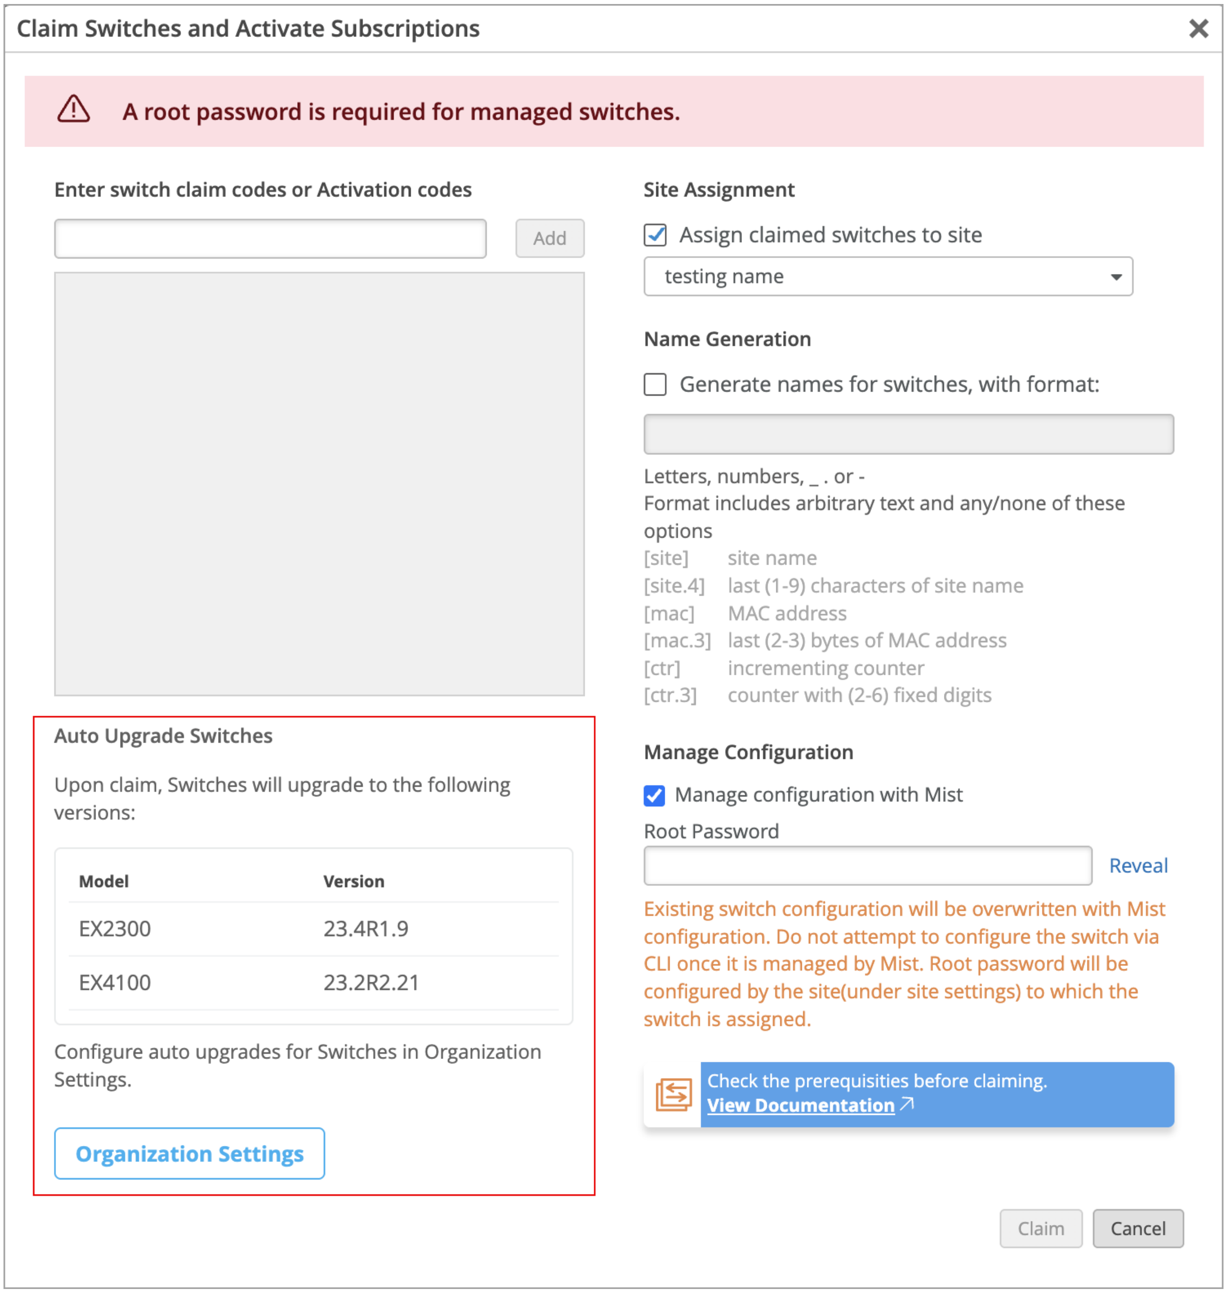

Claim Switches and Activate Subscriptions

A root password is required for managed switches.

Enter switch claim codes or Activation codes

Auto Upgrade Switches

Upon claim, Switches will upgrade to the following versions:

Model	Version
EX2300	23.4R1.9
EX4100	23.2R2.21

Configure auto upgrades for Switches in Organization Settings.

Site Assignment

☑ Assign claimed switches to site

testing name

Name Generation

□ Generate names for switches, with format:

Letters, numbers, _ . or -
Format includes arbitrary text and any/none of these options

[site]	site name
[site.4]	last (1-9) characters of site name
[mac]	MAC address
[mac.3]	last (2-3) bytes of MAC address
[ctr]	incrementing counter
[ctr.3]	counter with (2-6) fixed digits

Manage Configuration

☑ Manage configuration with Mist

Root Password

Existing switch configuration will be overwritten with Mist configuration. Do not attempt to configure the switch via CLI once it is managed by Mist. Root password will be configured by the site(under site settings) to which the switch is assigned.

Check the prerequisities before claiming.
View Documentation

Create Recovery Snapshot for a Switch

SUMMARY

To ensure that you have a full backup if you ever need it, follow these steps to create recovery snapshots for your switches.

A recovery snapshot stored in OAM (Operations, Administration, and Maintenance) volume holds a full backup that can be used in case something goes wrong with a Junos volume.

To create a recovery snapshot

1. Click **Switches** to go to the list of switches.
2. Locate the switch for which you want to create the snapshot and then click it to open the switch details page.
3. On the switch details page, click **Utilities** > **Snapshot Device**.

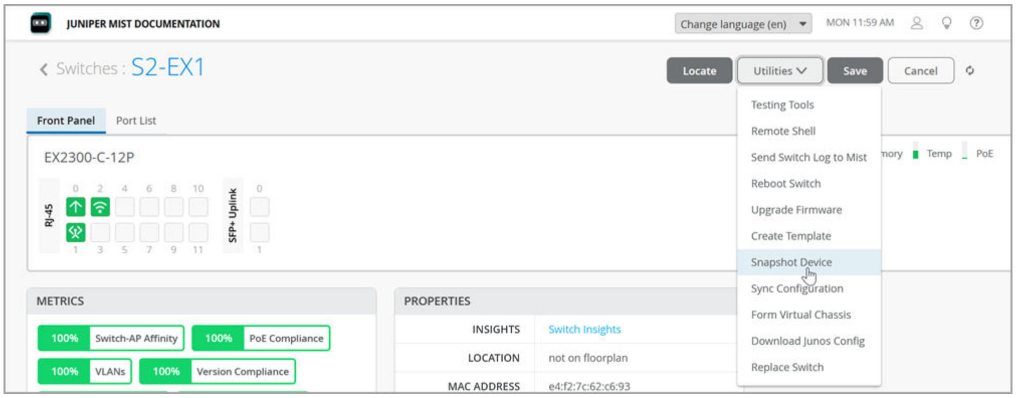

The **Snapshot Device** confirmation window appears.

4. On the **Snapshot Device** confirmation window, click **Request Snapshot.**

Rollback Junos to a Previous Version

SUMMARY

Follow these steps to revert to a previous version of Junos.

For EX4000 Series, EX4100 Series, and EX4400 Series switches, you can revert the version of Junos running on a standalone switch or Virtual Chassis device to the previous version stored on the device. Both the current and previous Junos versions must support Cloud X. In addition, the following conditions apply:

- All devices selected for the rollback operation, whether standalone switches or members in a Virtual Chassis device, must support rollback.
- Each device selected for the rollback operation must have a backup Junos version that matches the others.

To rollback to a previous version of Junos:

1. Click **Switches** to go to the list of switches.
2. Locate the switch for which you want to create the snapshot and then click it to open the switch details page.
3. On the switch details page, click **Utilities** > **Upgrade Firmware**.

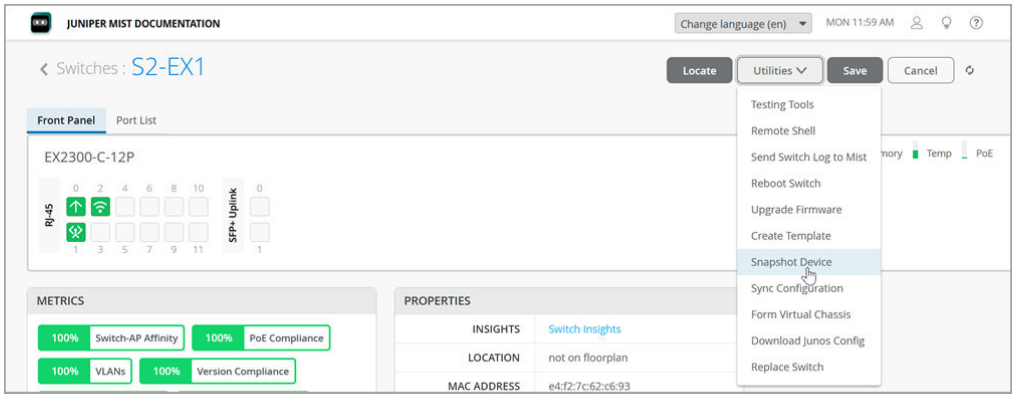

INSIGHTS Switch Insights
LOCATION not on floorplan
MAC ADDRESS e4:f2:7c:62:c6:93

The **Upgrade Switch Firmware** window appears.

4. On the **Upgrade Switch Firmware** page, click **Revert to Previous Running Version**.
5. Click **Start** to begin the rollback.

Assign a Role to Switches

SUMMARY

Assign a role to a switch so that it can be managed by role-based configuration rules in your configuration template.

You can select and apply a role to an individual switch from the switch details page or from the switch list. Use this feature to ensure that a switch does not inherit a role that does not exist. A switch can have only one role at a time.

You can create new switch roles as part of Select Switches Configuration rules in a switch template (**Organization** > **Switch Templates** or **Sites** > **Switch Configuration**). See also: "Configure Switches Using Templates" on page 21.

To assign a role to a switch:

1. Click **Switches** to go to the switch list.
2. Identify the switch that you want to assign a role to and then select the switch check box.
3. On the **More** menu, click **Assign Switch Role**.

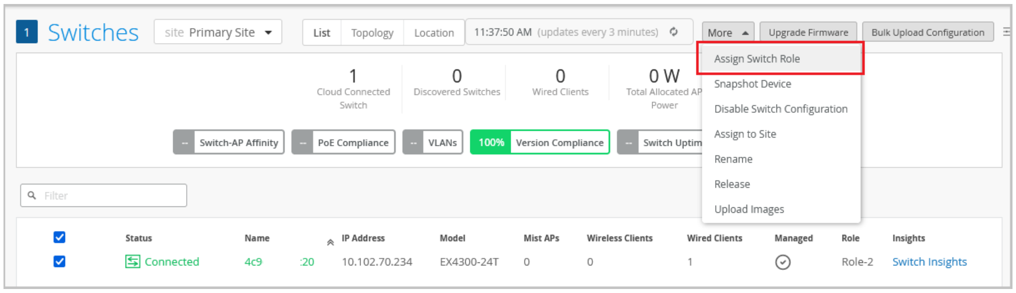

☑	Status	Name		IP Address	Model	Mist APs	Wireless Clients	Wired Clients	Managed	Role	Insights
☑	Connected	4c9	:20	10.102.70.234	EX4300-24T	0	0	1	✓	Role-2	Switch Insights

4. In the **Role** field, specify the switch role and then click **OK**. Click inside the Role field to view all the available roles. The Role field displays the switch roles that exist in the configuration template associated with the selected switch. Make sure that the role you assign to a switch exists in the configuration template associated with the selected switch.

   Alternatively, navigate to the switch details page (**Switches** > **Switch-Name**), specify the switch role in the INFO section, and save the configuration (see the image below).

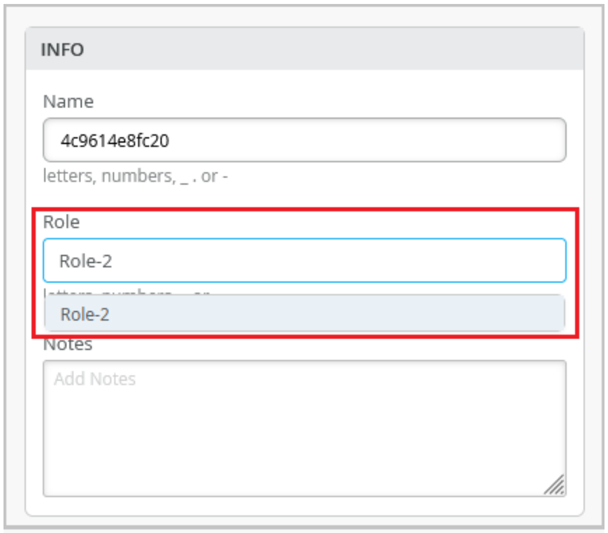

INFO

Name

4c9614e8fc20

letters, numbers, _ . or -

Role

Role-2

Role-2

Notes

Locate a Switch by LED

SUMMARY

Follow these steps to quickly find a particular switch or Virtual Chassis member by causing the status LED to flash.

To find a switch or Virtual Chassis member on a rack full of switches, you can enable the **Locate** beacon to illuminate an LED on the switch. For Virtual Chassis, this means you can **Locate** the primary, backup, or a given linecard member (only one member can be located at a time).

To turn on the LED beacon to identify the switch,

1. In the main menu, select **Switches** to open the list of switches.
2. Choose the switch you want to locate by clicking its name in the list of switches.

   The switch details page appears.
3. Click the **Locate** button at the top of the page to start the beacon, and again when you want to stop it.

   To locate a member of Virtual Chassis, select the member in the drop down that appears.

Figure 5: The Mist LED on an EX Series Chassis

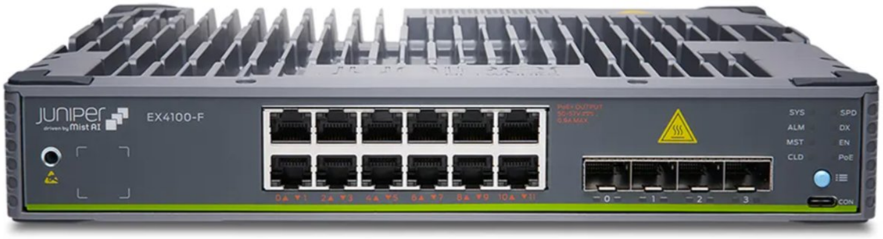

Replace a Switch

SUMMARY

Follow these steps to replace a switch from the Juniper Mist™ portal, without disrupting network services.

Before replacing a switch, ensure the following:

- The old switch that needs to be replaced is claimed or adopted by your organization and is assigned to a site. This switch can be in Connected or Disconnected state.
- The new switch being added is not assigned to any site in the organization. Furthermore, the new switch is listed on the Inventory page with the status Unassigned.

To replace a switch:

1. Click **Switches** on the left navigation pane of the Mist portal.
2. On the list tab, click the switch that needs to be replaced.
   The "switch details" on page 294 page appears.
3. Select the **Replace Switch** option from the **Utilities** drop-down list on the details.

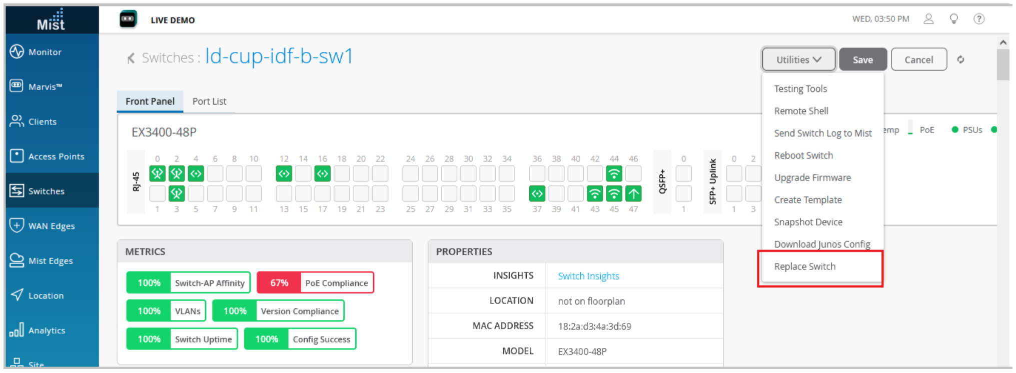

METRICS

100% Switch-AP Affinity	67% PoE Compliance
100% VLANs	100% Version Compliance
100% Switch Uptime	100% Config Success

PROPERTIES

INSIGHTS	Switch Insights
LOCATION	not on floorplan
MAC ADDRESS	18:2a:d3:4a:3d:69
MODEL	EX3400-48P

4. From the **MAC Address of the unassigned switch** drop-down list, select the MAC address of the unassigned switch. This is the replacement switch—the switch you will use to replace the switch that you selected in Step 2. If the Mist network doesn't include any unassigned switches, the Mist portal doesn't display unassigned switches. In that case, this drop-down list doesn't show any MAC addresses.

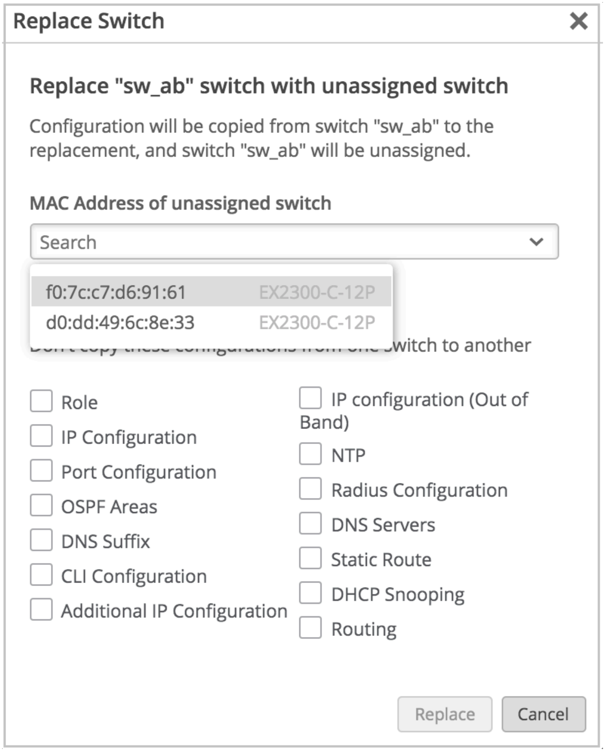

Replace Switch

Replace "sw_ab" switch with unassigned switch

Configuration will be copied from switch "sw_ab" to the replacement, and switch "sw_ab" will be unassigned.

MAC Address of unassigned switch

f0:7c:c7:d6:91:61	EX2300-C-12P
d0:dd:49:6c:8e:33	EX2300-C-12P

Don't copy these configurations from one switch to another

- □ Role
- □ IP Configuration
- □ Port Configuration
- □ OSPF Areas
- □ DNS Suffix
- □ CLI Configuration
- □ Additional IP Configuration
- □ IP configuration (Out of Band)
- □ NTP
- □ Radius Configuration
- □ DNS Servers
- □ Static Route
- □ DHCP Snooping
- □ Routing

By default, Mist copies the configuration of the existing switch to the new switch. To discard any specific configurations that you don't want to copy to the new switch, select the appropriate check boxes on the Replace Switch window.

If the new switch has a different number of ports than the switch being replaced, Mist discards the port configuration automatically. If the current switch template on the site doesn't cover the configuration requirement of your new switch, we recommend that you assign a different template. Assign your site a template with a configuration that meets the requirements of the new switch. See "Configure Switches Using Templates" on page 21.

5. Click **Replace**.

   When Mist replaces the switch, the new switch takes the place of the old switch on the Mist portal. The status of the old switch changes to Unassigned. You can view the old switch on the Inventory page.

Switch replacement using APIs

To replace a switch using APIs, make a POST API call as shown in the example below:

```
POST /api/v1/orgs/:org_id/inventory/replace

{
   "site_id": "4ac1dcf4-9d8b-7211-65c4-057819f0862b",
   "mac": "5c5b35000101",
   "inventory_mac": "5c5b35000301",
   "discard": []
}
```

On the **discard** list, you can specify the attributes that you do not want to copy to the new switch configuration. If the **discard** list is blank, Mist copies all the existing switch attributes from the old switch to the new switch.

NOTE:

- If a switch with a higher number of ports is being replaced with a switch with a lower number of ports, the port configuration is applied only to the ports with overlapping port numbers. The rest of the port configurations are discarded.
- If a switch with mge ports is being replaced with a switch with ge ports or vice versa, the port configurations are not applied to the switch.

Remote Shell Access

SUMMARY

Configure remote access from the Mist console to the Junos shell for connected devices.

By default, super users and network administrators can open a remote shell connection to any connected EX, SRX, and QFX device,managed or unmanaged, from the Mist console. This is especially useful for running Junos operational (show) commands for troubleshooting, viewing the configuration, or checking connection details and statistics. Note that you can also enter Junos configuration mode to edit the CLI directly, although this practice is discouraged for managed devices because those changes are not visible from Mist console, and any such changes will be eventually overwritten. Instead, we recommend using the CLI configuration for adding Junos commands that are not available in the Mist console.

Super users and network administrators can globally disable remote shell access. In other words, for all devices in the organization, and all Mist account levels, new remote connections from the Mist console will be denied. Any existing remote shell connections will remain active, though.

Super users and network administrators can also allow read-only shell access (`monitor` , `show`, `test`, `quit`, `help`, `request session`, `ssh user@localhost`), for the following Mist account types:

- Helpdesk
- Switch Port Operator
- Observer
- Super Observer

Configure remote shell access from the **Device Management** section of the **Organization > Settings** page.

Figure 6: Remote Shell Access

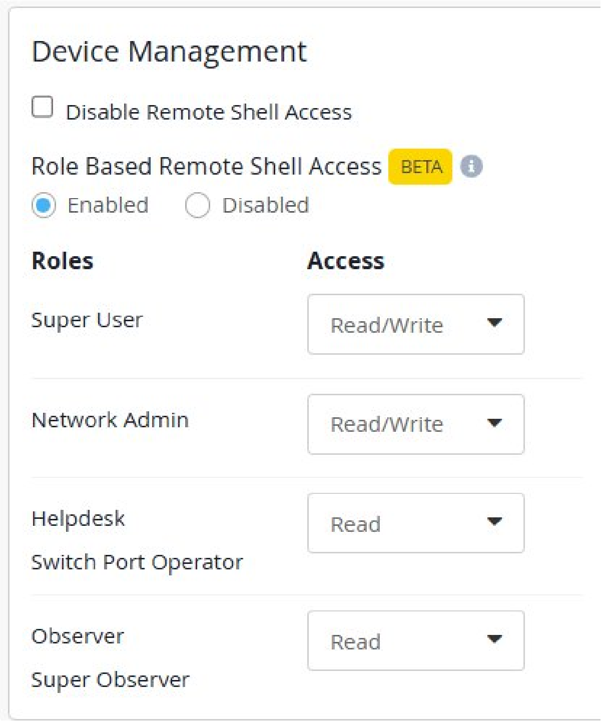

Device Management

☐ Disable Remote Shell Access

Role Based Remote Shell Access BETA

● Enabled ○ Disabled

Roles	Access
Super User	Read/Write
Network Admin	Read/Write
Helpdesk Switch Port Operator	Read
Observer Super Observer	Read

When you open a remote shell to a device from the Mist console, the Junos privileges will reflect those of your Mist account, either *mist-web-admin* or *mist-web-viewer*. Configuration changes are shown on the Insight page as, commit user: *mist-web-admin*.

Mist will use the existing **lo0** (loopback) interface for the connection, if it exists. Otherwise, Mist will provision the interface with an IP address of 127.127.127.1/32, and create a firewall rule, if needed, to allow the traffic.

NOTE: Open a remote shell to a device by clicking **Utilities** > **Remote Shell** from the configuration page. For example, in the Switches list view page, click a switch name and then look for Utilities in the upper right corner of the page that opens.

Connect a Switch to Mist Cloud via a Proxy Server Using Cloudx

SUMMARY

Follow this workflow to connect an EX Series switch to the Mist cloud by using a proxy server.

IN THIS SECTION

Juniper CloudX Overview

Juniper CloudX, integrated natively into Junos OS, is an advanced architecture that ensures faster and secure communication between Juniper switches and the Mist cloud. It is responsible for creating a secure connection between the switch and the Mist cloud. CloudX-enabled switches can be monitored and managed by cloud services.

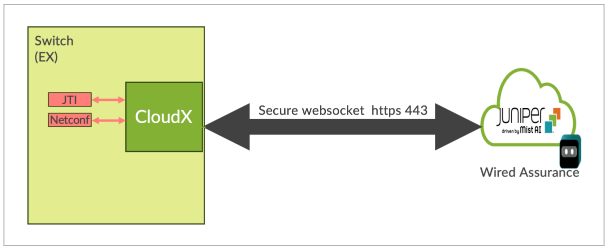

CloudX applies to both new and existing switches. It enables the new switches to communicate directly over HTTPS 443 when they are onboarded to Mist cloud via ZTP. With CloudX enabled, the existing switches that are connected to the Mist cloud via TCP port 2200 will have their connection switched to CloudX with no impact on the data plane. For switches to connect and communicate using CloudX over TCP 443, the following firewall port must be opened: jma-terminator.[xx].mist.com(TCP 443). The variable [xx] should be replaced by the environment name.

Benefits of CloudX:

- Keeps the data on the cloud up to date. Events are sent to the cloud every 10-15 seconds and stats are updated every 60 seconds.
- Leverages the Junos Telemetry Interface (JTI), which ensures asynchronous and faster communication by bypassing any polling from the cloud to the switch.
- Enables switches to connect to the cloud over HTTPS port 443, like Mist APs. You do not need to open any non-standard ports on the firewall.
- Enables switches to communicate with the Mist cloud via a proxy server. You can statically define a proxy server or dynamically send proxy server details via DHCP Option 43. For more information, see "Connect a Switch to Mist Cloud via a Proxy Server Using Cloudx" on page 181.
- Offers packet capture for switches on the Mist Cloud. You can initiate packet capture on a single switch port or a range of ports. You can leverage the on-demand packet capture feature in Mist to view transit traffic or control traffic. For more information, refer to "Enable Packet Capture on a Switch" on page 186.

Availability of CloudX

The following table lists the platforms that support CloudX in different Junos releases. The table lists multiple Junos versions for each platform. Different models (variants) within each platform are also supported. So, the EX4100-F variant of the EX4100 Series is also supported. We recommend that you upgrade the switch to a Junos suggested release for the CloudX support.

> **NOTE**:
> For CloudX to work, you must ensure that the firewall port towards jma-terminator.xx.mist.com is open and SSL decryption is disabled on the firewall (for more information, refer to Juniper Mist Firewall Ports and IP Addresses for Firewall Configuration). To check if your switch is communicating with Mist cloud by using CloudX, refer to the steps listed in Troubleshooting Juniper CloudX. If you still don't see CloudX enabled on your switch even after upgrading it to a supported Junos release, contact Juniper support.

Table 10: CloudX-Supported Platforms

Platforms	Supported Junos Release	**CloudX Availability**
EX2300/EX3400	23.4R2-S4 and above 24.2R1-S2 and above	Generally Available

Table 10: CloudX-Supported Platforms *(Continued)*

Platforms	Supported Junos Release	CloudX Availability
EX4000	24.4R1 and above 24.4R1-S2 and above	Generally Available
EX4400/EX4100	22.4R2-S1 and above 22.4R3 and above 23.4R2 and above 24.2R1 and above	Generally Available
EX4650/QFX5120	23.4R2-S4 and above 24.2R1-S2 and above	Generally Available

NOTE:

- The following Junos versions do not support CloudX: 23.1R1 and 22.1-22.3.
- All variants of each switch model listed in this table also support CloudX. For example, if the table lists EX4100, assume that the EX4100-F switches also support CloudX.

If you face any issues with CloudX, you can troubleshoot it by following the steps listed in "Troubleshooting Juniper CloudX" on page 332.

Connect a Switch to Mist Cloud via a Dynamic Proxy Server

Before you connect the switch to the Mist cloud via a dynamic proxy server, ensure that the following prerequisites are met:

- The switch is onboarded to the Mist cloud using the claim code or activation code.
- The switch is running a CloudX-supported Junos version. For more information, see Juniper CloudX Overview.
- The DHCP server is able to hand out the proxy server information (via Option 43) and other elements such as IP Address, DNS, and default route.

- The switch can reach the HTTP proxy server over an IP network.
- The HTTP proxy server can redirect traffic to the Mist cloud. This example shows how to configure the proxy server:

To connect a switch to Mist cloud via a dynamic proxy server:

1. Power on the switch.
2. Connect the switch to the uplink (via OOB or in-band port).

   The switch sends a DHCP Discover message and accepts the Offer message along with DHCP proxy server information sent via Option 43. The switch stores the proxy server information at `/var/etc/phc_vendor_specific_info.xml`. The switch reaches out to the proxy server during the ZTP boot-up process and connects to the Mist cloud via HTTP proxy server.
3. Log in to the switch and verify the connectivity to the Mist cloud by using the following CLI command:

   ```
   show system connections | grep port used for connectivity between switch and proxy offered by DHCP
   ```
4. In case the switch does not connect to the cloud, collect logs from the following files on the switch and create a support ticket:

   /var/log/mcd.log, /var/log/messages and RSI

   For more information, see How to collect logs and files from standalone and Virtual Chassis/VCF devices.

Connect a Switch to Mist Cloud via a Static Proxy Server

If a switch cannot receive the proxy information via DHCP, you can configure it with a static proxy server through which the switch can connect to the Mist cloud. In this case, the DHCP server does not hand out the proxy server information via Option 43.

Before you connect the switch to the Mist cloud via a static proxy server, ensure that the following prerequisites are met:

- The switch is onboarded to the Mist cloud using the claim code or activation code.
- The switch has the configuration management option enabled in Mist. If not, you will need to use the switch CLI to configure the proxy server.
- The switch is running the Junos version 21.4R3-S4, 22.4R2-S1, or above.
- The local DHCP server is able to hand out IP address, DNS, default route, or statically defined route on the switch. This process involves staging the switch before establishing the cloud connectivity. If this prerequisite is met, the switch will be able to reach the HTTP proxy server over an IP network.

To connect a switch to Mist cloud via a static proxy server:

1. Log in to the Mist portal (manage.mist.com).
2. Click **Organization** > **Site Configuration** > ***site-name*** to navigate to the site where the switch is onboarded.
3. On the Site Proxy tile of the site configuration page, configure the proxy information, as shown below:

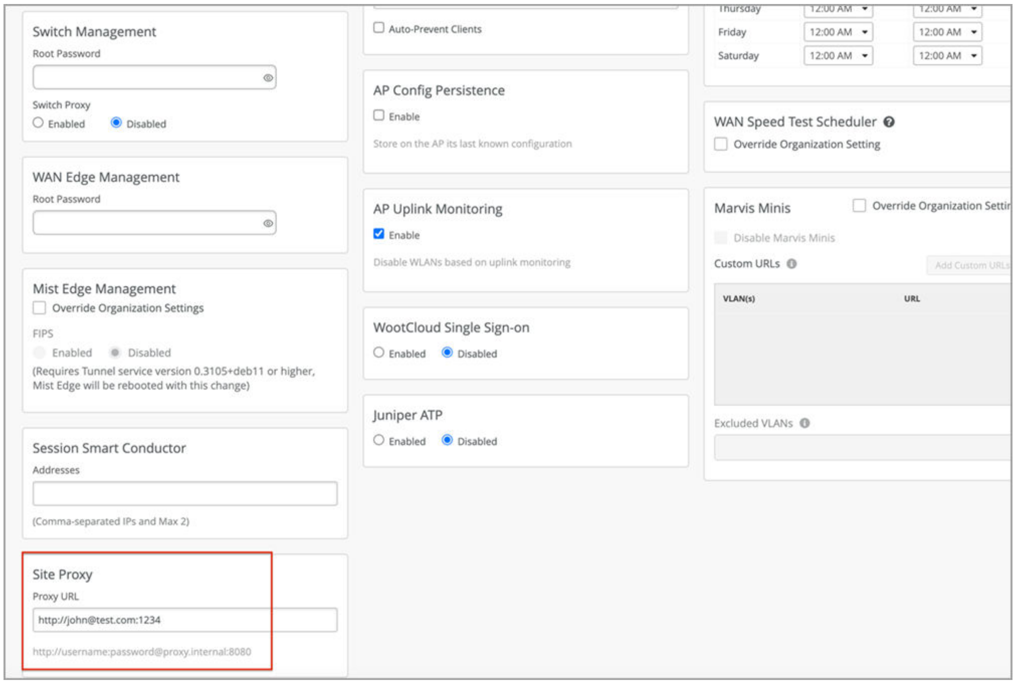

Switch Management
Root Password
Switch Proxy ○ Enabled ● Disabled

WAN Edge Management
Root Password

Mist Edge Management
□ Override Organization Settings
FIPS ○ Enabled ● Disabled
(Requires Tunnel service version 0.3105+deb11 or higher, Mist Edge will be rebooted with this change)

Session Smart Conductor
Addresses
(Comma-separated IPs and Max 2)

Site Proxy
Proxy URL: http://john@test.com:1234
http://username:password@proxy.internal:8080

□ Auto-Prevent Clients

AP Config Persistence
□ Enable
Store on the AP its last known configuration

AP Uplink Monitoring
☑ Enable
Disable WLANs based on uplink monitoring

WootCloud Single Sign-on ○ Enabled ● Disabled

Juniper ATP ○ Enabled ● Disabled

Thursday	12:00 AM	12:00 AM
Friday	12:00 AM	12:00 AM
Saturday	12:00 AM	12:00 AM

WAN Speed Test Scheduler
□ Override Organization Setting

Marvis Minis □ Override Organization Settir
□ Disable Marvis Minis
Custom URLs

VLAN(s)	URL

Excluded VLANs

4. Stage the switch in a non-proxy environment to connect to the cloud. If the switches are staged, they can gather the proxy information from the Mist cloud (which you configured in the previous step). Staging in this context means connecting the switch to the Mist cloud in a non-proxy environment before deployment.

   When you complete the above steps, the switch will be able to reach the proxy server during the ZTP process.
5. If the switch is not connecting to the proxy server during the ZTP process, flap the uplink port to force the switch to connect to the proxy server.
6. Log in to the switch and verify the connectivity to Mist cloud by using the following CLI command:

   ```
   show system connections | grep port used for connectivity between switch and proxy
   ```
7. In case the switch is not connecting to the cloud, collect logs from the following files on the switch and open a support case with Juniper support:

   /var/log/mcd.log, /var/log/messages and RSI

   For more information, see How to collect logs and files from standalone and Virtual Chassis/VCF devices.

Enable Packet Capture on a Switch

Packet capture (PCAP) is a tool that helps you to analyze network traffic and troubleshoot network problems. The packet capture tool captures real-time data packets traveling over the network for monitoring and logging.

You can enable on-demand PCAP on switches that have CloudX running. For the list of switches that support CloudX, refer to Juniper CloudX Overview.

The PCAP feature captures both control traffic (the traffic handled by the device CPU) and transit traffic (the traffic forwarded by network processors) that pass through switches at a site.

> **NOTE**: To capture transit traffic, a switch must have the secure PCAP feature enabled. Currently, only the EX4400, EX4100, EX4000 switches support this feature. PCAP captures only ingress transit traffic, not egress traffic.

Packets are captured as binary data, without modification. You can read the packet information offline with a packet analyzer such as Wireshark or tcpdump.

To enable PCAP on a switch:

1. Select **Site** > **Wired** > **Switch Packet Captures**.
2. Select a site from the **Site** drop-down list.
3. Click the + icon next to the **Add Switch** field, and select the switch on which you want to enable packet capture.

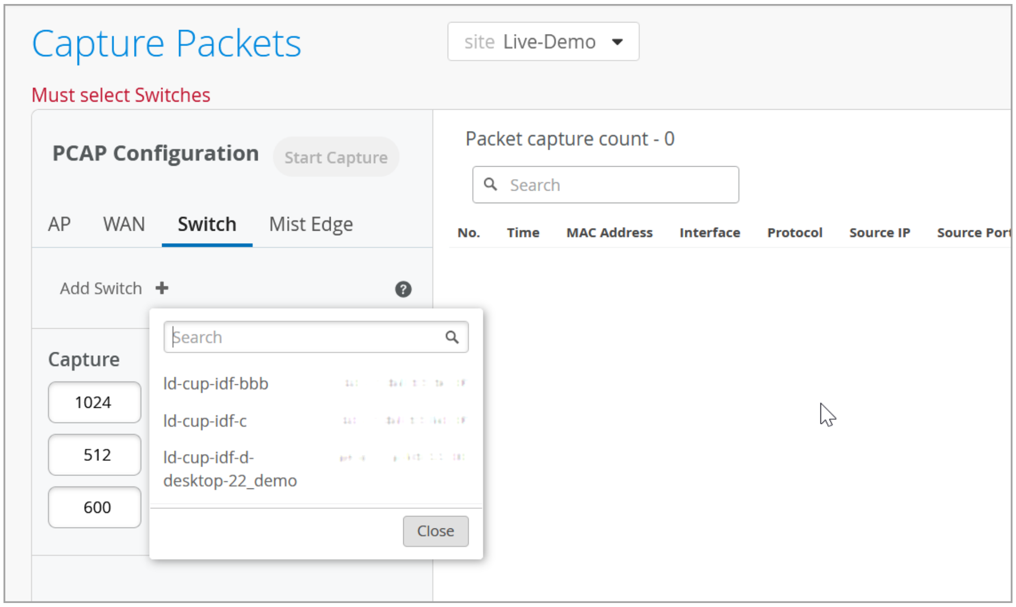

You can select multiple switches for a single packet capture operation.

4. Specify the number of packets captured per switch, packet size in bytes, and the duration of the capture session in seconds.

> **NOTE**: If you specify 0 in the **No. of packets/Switch** field, unlimited number of packets will be captured.

5. Configure a port filter for packet capture. To do that follow the steps below:

   a. Click **Add Port Filter**.

   b. Click the port icon in the **Port Name** field, select a port on which you want to enable packet capture, and then click **Done**.

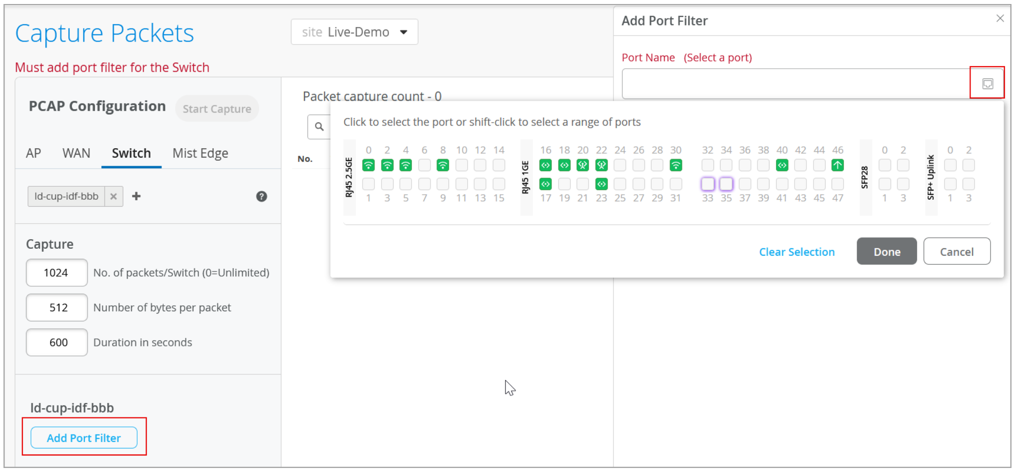

Capture Packets

site Live-Demo

Must add port filter for the Switch

PCAP Configuration

Capture

1024 No. of packets/Switch (0=Unlimited)

512 Number of bytes per packet

600 Duration in seconds

ld-cup-idf-bbb

Add Port Filter

Port Name (Select a port)

Click to select the port or shift-click to select a range of ports

> **NOTE**: You can select multiple ports from multiple switches in a single packet capture configuration.

   If you want to capture traffic on CPU, select the **Capture Traffic on CPU** check box.

   c. Under Advanced filters, specify filters using a tcpdump expression if required.

   You can also use the expression builder to build the expression.

   d. Click **Save**.

6. Click **Start Capture** to enable packet capture on the selected port.

After the packet capture is complete, you can download the file for inspection. To do that, click **Captured File** on the upper right of the screen.

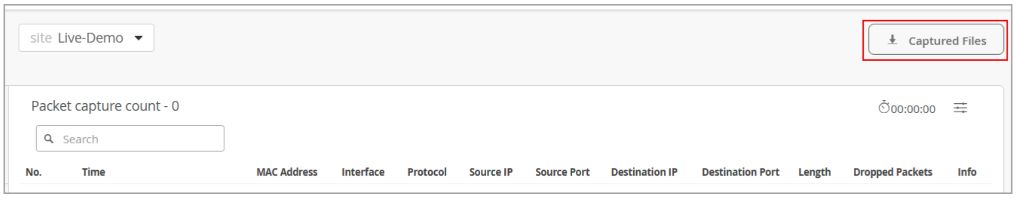

site Live-Demo

Captured Files

Packet capture count - 0

00:00:00

Search

No.	Time	MAC Address	Interface	Protocol	Source IP	Source Port	Destination IP	Destination Port	Length	Dropped Packets	Info

To know more about how to view the packet capture in Wireshark, refer to Configure IEEE 802.11 on Wireshark and View Wireless Packet Captures in Wireshark.

Configure the System Log

SUMMARY

Send system log messages to files, remote destinations, user terminals, or the system console.

IN THIS SECTION

Junos OS generates system log messages (also called *syslog messages*) to record events that occur on a switch, including the following events:

- Routine operations, such as creation of an Open Shortest Path First (*OSPF*) protocol adjacency or a user login to the configuration database
- Failure and error conditions, such as failure to access a configuration file or unexpected closure of a connection to a peer process
- Emergency or critical conditions, such as shutdowns due to excessive temperature

For the switches that you manage in the Juniper Mist portal, you can configure syslog in the switch settings.

To configure the system log:

1. Go to the switch template or the site-level configuration:
   - To configure an organization-level switch template—From the left menu, select **Organization** > **Wired** > **Switch Templates**.
   - To configure site-level switch settings—From the left menu, select **Site** > **Wired** > **Switch Configuration**. Click the site that you want to configure.
2. Under **All Switches Configuration**, find the **Syslog** section.
3. If you're on the site-level configuration page, select **Override Configuration Template**.

4. Click **Enabled**.

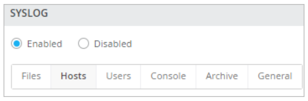

SYSLOG

● Enabled ○ Disabled

NOTE: After you enable logging, you must set up at least one type of logging.

5. Use the tabs to configure the log settings.

   For help with these options, see "Syslog Options" on page 189.

6. Click **Save** at the top-right corner of the configuration page.

Syslog Options

SUMMARY

Use this information to send system log messages to local files, external hosts, users, or the console. Also configure content types, archive settings, and more. (Available logging options vary, depending on which settings you're configuring, such as syslog for a switch, a WAN Edge device, or other.)

Table 11: Files

Field	Description
Add File	Use the **Add File** link to specify a named file in the local file system. Syslog messages that match the specified criteria will be sent to this file. Cycle through this process as needed to specifying multiple files for different matching criteria.
File Name	Enter a name for the file.
Match	To filter messages, enter a text string. Only log messages that contain this string will be saved to the specified file.

Table 11: Files *(Continued)*

Field	Description
Explicit Priority	Select this option to include priority information in the file.
Structured Data	Select this option to use structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.
Archive	Set the maximum number of files to store, and the maximum file size in megabytes (m) or bytes.
Contents	To specify the types of log information to capture, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be saved to the specified file. For more information about the content options, see Table 14 on page 192 later in this topic.

Table 12: Hosts

Field	Description
Add Host	Use the **Add Host** link to specify an external server that is configured as a system log message host. Syslog messages that match the specified criteria will be sent to this server. Cycle through this process as needed to specifying multiple hosts for different matching criteria.
Host	Enter the host.
Port	Enter the port number, from 1 through 65535.

Table 12: Hosts *(Continued)*

Field	Description
Match	To filter messages, enter a text string. Only log messages that contain this string will be saved to the specified file.
Explicit Priority	Select this option to include priority information in the file.
Structured Data	Select this option to use structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.
Archive	Set the maximum number of files to store, and the maximum file size in megabytes (m) or bytes.
Contents	To specify the types of log information to capture, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be saved to the specified file. For more information about the content options, see Table 14 on page 192 later in this topic.

Table 13: Users

Field	Description
Add User (link)	On the Users tab, you're sending log messages to the terminal session of one or more users. To add a user, click this link, and the enter the user information. Cycle through this process as needed to specify multiple users for different matching criteria.
User	Enter one or more usernames, separated by spaces, or enter an asterisk (*) to include all users.

Table 13: Users *(Continued)*

Field	Description
Match	To filter messages, enter a text string. Only log messages that contain this string will be sent to the specified users.
Contents	To specify the log information to capture for this user, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be sent to the specified users. For more information about the content options, see Table 14 on page 192 later in this topic.

Table 14: Console Options and Content Types

Field	Description
Add Content (link)	Use the **Add Content** link to specify the types of log messages to capture. You'll cycle through this process multiple times to add different content types. For example, let's say you want to capture (1) critical authorization events, (2) warnings for authorization events, (3) critical change logs, (4) change log errors, and (5) ftp errors. For this example, you'd create five content types. **NOTE**: You can add content types in a few different places. For example, if you're working on the Console tab, you're specifying the content to send to the console. If you're working on the User tab, you're specifying the content to capture for that user.
Facility	Select a log event that you want to capture in this content type.
Level	For the specified log event, select the severity level to capture in this content type.

Table 15: Archive Tab

Field	Description
Files	Specify the maximum number of files (1-1000) to store in the system log. After this maximum is reached, log messages are archived.
Size	Specify the maximum size (65536 - 1073741824) of the system log. Select either **m** (megabytes) or **bytes**. After this maximum is reached, log messages are archived.

Table 16: General Tab

Field	Description
Time Format	Use this option to include the milliseconds, the year, or both in the syslog timestamps. • To add an option, click the plus (+) button, then click the option in the list. • To remove an option, click the X.
Routing Instance	By default, system logging traffic is sent from the management interface on your device and its associated routing instance. However, the Junos syslog client is completely VRF aware. If a server is reachable through a virtual routing and forwarding (VRF) instance, the syslog client can send log messages to the server. Use this option to specify a routing instance to use.
Network	Use this option to specify a network instance to use.

Release a Switch from Inventory

SUMMARY

Follow these steps to remove an unneeded switch from your organizations inventory in Juniper Mist™.

If you no longer want to include a switch in your Juniper Mist organization, you can release it from your inventory.

> **NOTE**: If the switch being released is part of any campus fabric topology, you must remove the switch from that topology before proceeding to release the switch from the inventory. To do that, follow the steps below:
>
> 1. From the campus fabric page (**Organization** > **Wired** > **Campus Fabric**), select the site to which the switch belongs.
> 2. Click the campus fabric topology to which the switch is linked.
> 3. Click **Edit Configuration** > **Continue** to navigate to the **Nodes** tab.
> 4. Select the switch to be removed and then click the **Remove From Topology** button at the lower right of the page.
> 5. Navigate to the **Confirm** tab by using the **Continue** button and then click **Apply Changes**.

1. From the left navigation menu of the Juniper Mist portal, select **Organization** > **Inventory**.
2. Click the **Switches** button at the top of the page.
3. Select the check box for one or more switches.
4. Click the **More** button near the top-right corner of the page, and then click **Release**.

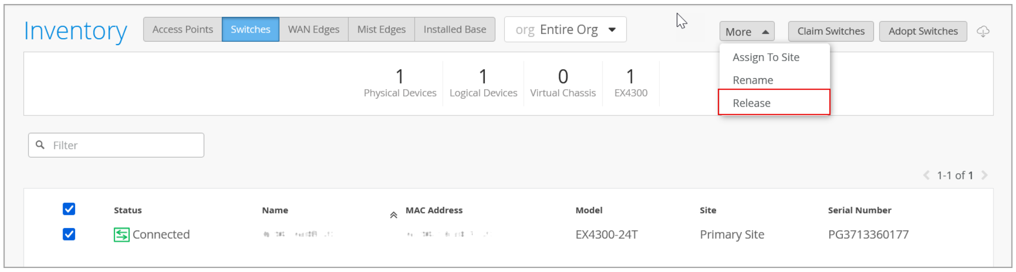

1	1	0	1
Physical Devices	Logical Devices	Virtual Chassis	EX4300

☑	Status	Name	MAC Address	Model	Site	Serial Number
☑	Connected	[illegible]	[illegible]	EX4300-24T	Primary Site	PG3713360177

5. When the confirmation window appears, click **Yes** to confirm that you want to release this switch.

   The switch is no longer claimed by this organization and no longer managed by Juniper Mist.

   You can release the switch from the Switches page (the switch list accessed by clicking **Switches** on the left navigation menu) as well.

Firewall Filters

SUMMARY

To ensure more granular control over network access, you can now set up access control lists (ACLs) for your Juniper Mist-managed switches. You can use port profile-based filters and RADIUS-based filters.

IN THIS SECTION

About Filters

You can create ACLs by using the following types of filters:

- **Port Profile-Based Filter**—This ACL will be applied as a layer 2 filter on the switchport in the input direction for all ports where the specified port profile is applied.
- **RADIUS-Based Filter**—This ACL uses a RADIUS based filter to filter traffic. The enforcement of each policy happens via the RADIUS server. These filters are supported on all EX Series switches that authenticate users through your RADIUS server. After adding your RADIUS firewall filters, make note of the IDs. You'll need them to create the switch policies in the Juniper Mist portal.

Create Firewall Filters

Firewall filters are configured in the form of policy labels and applied to switch policy rules. These labels are used to categorize and classify users (as sources) and resources (as destinations). Then you use these labels in your switch policies to specify which users are allowed to access specific resources within the network. You can define the labels at the organization, site, or switch level.

Before You Begin

- For port-profile based filters, first set up your port profiles.
- For RADIUS-based firewall filters, first set them up on your RADIUS server. Add the filters by using filter-id attribute in the Juniper dictionary on your RADIUS server. For help with RADIUS filters, see Configuring Firewall Filters on the RADIUS Server in the *User Access and Authentication Administration Guide for Junos OS*.

To create firewall filters:

1. In the Juniper Mist portal, navigate to your switch template or site-level switch configuration.
    - To find a switch template—From the left menu, select **Organization** > **Wired** > **Switch Templates**. Then click the template.
    - To find a site-level switch configuration—From the left menu, select **Site** > **Wired** > **Switch Configuration**. Then click the site.
2. To configure labels, scroll down to the **Switch Policy Labels** section, and add your source and destination tags as described below.
    - **Source**—Click **Add Source**, and then refer to the on-screen text to enter the information. Click **Add** to save the new source.

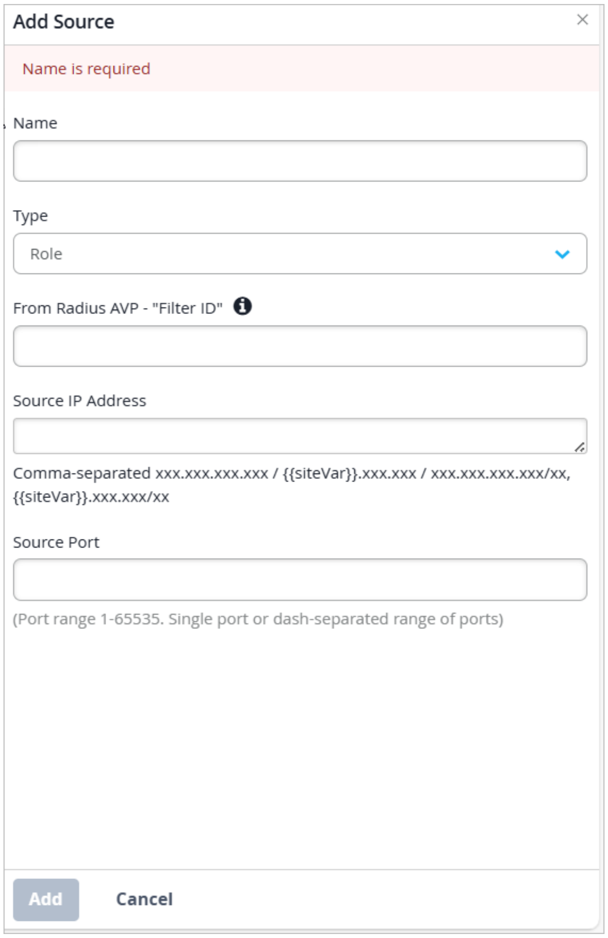

Add Source

Name is required

Name

Type

Role

From Radius AVP - "Filter ID"

Source IP Address

Comma-separated xxx.xxx.xxx.xxx / {{siteVar}}.xxx.xxx / xxx.xxx.xxx.xxx/xx, {{siteVar}}.xxx.xxx/xx

Source Port

(Port range 1-65535. Single port or dash-separated range of ports)

Table 17: Source Settings

Field	Description
Name	Enter a name to identify this source.
Type	Select the source type.
Port Profile	If you selected **Port Profile** as the **Type**, select a profile from this drop-down menu. This menu includes system-defined profiles and others that you've added to the configuration.
From Radius AVP - "Filter ID"	For a RADIUS-based filter, enter the filter. For more information, hover over the **i** button next to this field.

Table 17: Source Settings *(Continued)*

Field	Description
Source IP Address	(Optional) Enter a single IP address or multiple addresses separated by commas.
Source Port	(Optional) Enter a port number or a range of numbers. Valid numbers are 1 to 65535. For a range, enter the start number, a dash, and the end number, such as 50-60.

- **Destination**—Click **Add Destination**, and then refer to the on-screen text to enter the information. Click **Add** to save the new destination.

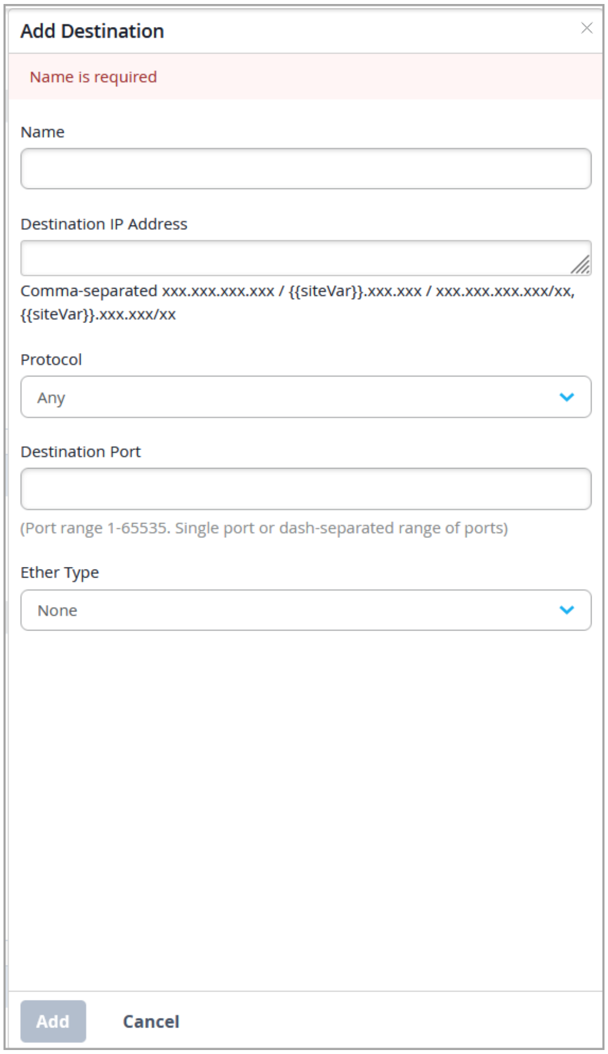

Add Destination

Name is required

Name

Destination IP Address

Comma-separated xxx.xxx.xxx.xxx / {{siteVar}}.xxx.xxx / xxx.xxx.xxx.xxx/xx, {{siteVar}}.xxx.xxx/xx

Protocol: Any

Destination Port

(Port range 1-65535. Single port or dash-separated range of ports)

Ether Type: None

3. To configure a policy, scroll down to the **Switch Policy** section and add a policy that uses your new labels:

   a. If you're working on a site-level switch configuration, select **Override Template Defined**.

   b. Click **Add Switch Policy**.
   A new policy appears at the top of the policy list, with a default name such as Switch Policy 1.

   c. Click the default policy name, and then enter a short, specific name to identify this policy.

   d. Under **Source**, click the **+** button, and then select a source label from the list.

   e. Under **Destination**, click the **+** button, and then select a destination label.
   The destination label appears with a green background, indicating that this policy allows access to this destination.

   f. If you want this policy to block access to the specified destination, click the destination button, and then click **Deny**.

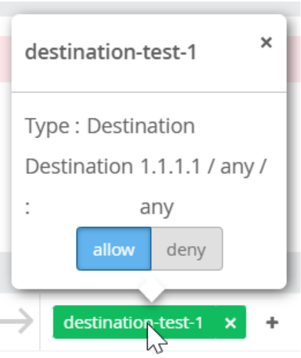

destination-test-1

Type : Destination

Destination 1.1.1.1 / any / : any

   When you change the policy to deny, the destination button turns red.

   NOTE:

   - The default action for a new rule is **allow** (green button).

   - Junous processes the policies in the listed order. After the final policy is processed, Junos defaults to *deny*. You don't need to add an explicit *deny* policy at the end of the list.

This example shows how your configuration might look after you add source labels, destination labels, and policies.

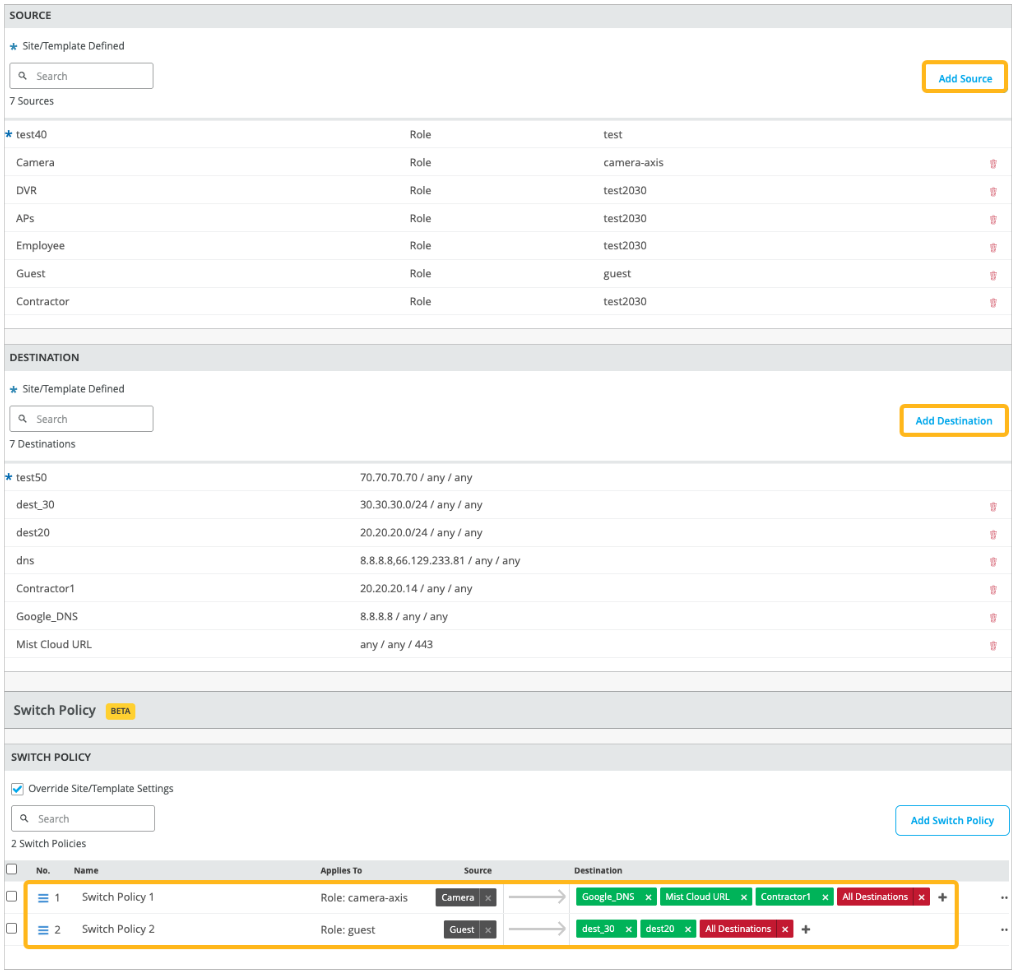

SOURCE

* Site/Template Defined

Search

Add Source

7 Sources

* test40	Role	test
Camera	Role	camera-axis
DVR	Role	test2030
APs	Role	test2030
Employee	Role	test2030
Guest	Role	guest
Contractor	Role	test2030

DESTINATION

* Site/Template Defined

Search

Add Destination

7 Destinations

* test50	70.70.70.70 / any / any
dest_30	30.30.30.0/24 / any / any
dest20	20.20.20.0/24 / any / any
dns	8.8.8.8,66.129.233.81 / any / any
Contractor1	20.20.20.14 / any / any
Google_DNS	8.8.8.8 / any / any
Mist Cloud URL	any / any / 443

Switch Policy BETA

SWITCH POLICY

☑ Override Site/Template Settings

Search

Add Switch Policy

2 Switch Policies

□	No.	Name	Applies To	Source		Destination
□	≡ 1	Switch Policy 1	Role: camera-axis	Camera ×	→	Google_DNS × Mist Cloud URL × Contractor1 × All Destinations × +
□	≡ 2	Switch Policy 2	Role: guest	Guest ×	→	dest_30 × dest20 × All Destinations × +

4. If you need to change the order of the policies, position your mouse on the blue button next to the policy number and drag up or down in the list.
5. Click **Save** at the top-right corner of the configuration page.
6. Review the on-screen confirmation information, and then click **Save** at the bottom of the confirmation window.

The switch policy configuration is pushed to the switches.

You can also find information about the number of times a switch policy rule was triggered (that is, matched by network traffic) at the switch level. The Switch Policy section on the switch details page provides the following details:

- Overall hit count for a switch policy. This information is displayed in the **Hit Count** column.

- Per-destination hit count for more granular insights. You can click each destination tag to view the hit count for that tag along with a policy trigger event time series.

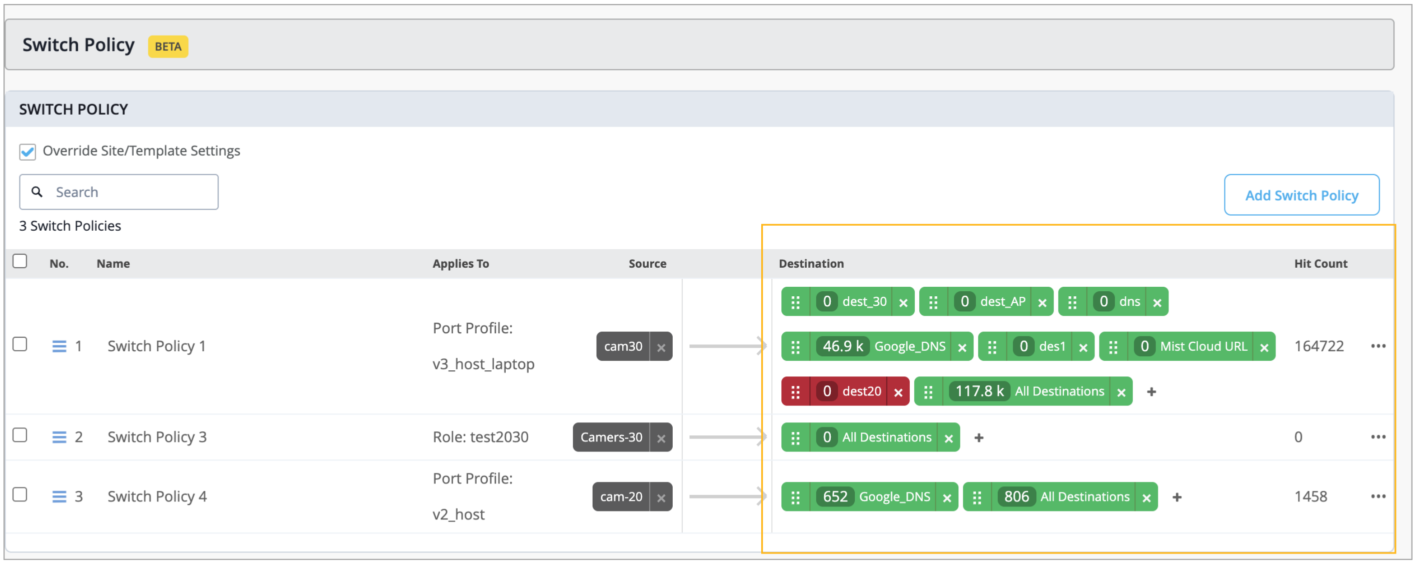

Switch Policy BETA

SWITCH POLICY

☑ Override Site/Template Settings

Search

Add Switch Policy

3 Switch Policies

□	No.	Name	Applies To	Source	Destination	Hit Count
□	1	Switch Policy 1	Port Profile: v3_host_laptop	cam30	0 dest_30 0 dest_AP 0 dns 46.9 k Google_DNS 0 des1 0 Mist Cloud URL 0 dest20 117.8 k All Destinations	164722
□	2	Switch Policy 3	Role: test2030	Camers-30	0 All Destinations	0
□	3	Switch Policy 4	Port Profile: v2_host	cam-20	652 Google_DNS 806 All Destinations	1458

Set Up Firewall Filters with Aruba ClearPass

> **NOTE**: As an Aruba ClearPass admin, use this procedure for general guidance only. For detailed and up-to-date information, see the Aruba ClearPass documentation and support site.

1. In Aruba ClearPass Policy Manager, navigate to **Configuration** > **Enforcement** > **Profiles**.
2. On the Enforcement Profiles page, add a profile, or edit an existing one.
   - Profile tab—Select **Filter ID Based Enforcement** as the template. Specify the required parameters.

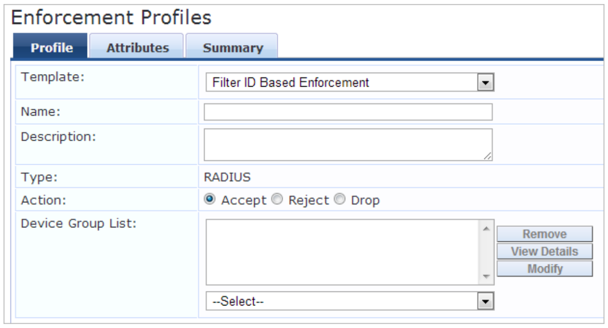

Enforcement Profiles

Profile | Attributes | Summary

Template:	Filter ID Based Enforcement
Name:	
Description:	
Type:	RADIUS
Action:	● Accept ○ Reject ○ Drop
Device Group List:	--Select--

   - **Attributes** tab—Add your RADIUS attributes.

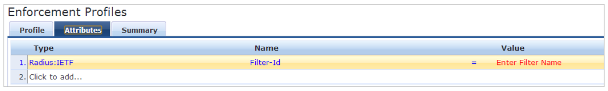

Enforcement Profiles

Profile | Attributes | Summary

	Type	Name		Value
1.	Radius:IETF	Filter-Id	=	Enter Filter Name
2.	Click to add...			

For help with the parameters, see your Aruba ClearPass documentation, such as Filter ID Based Enforcement Profile.

Set Up with Cisco ISE

NOTE: As an Cisco ISE admin, use this procedure for general guidance only. For detailed and up-to-date information, see the Cisco ISE documentation and support site.

In Cisco ISE, navigate to **Policy** > **Policy Elements** > **Results** > **Authorization** > **Authorization Profiles**. Enter the required parameters.

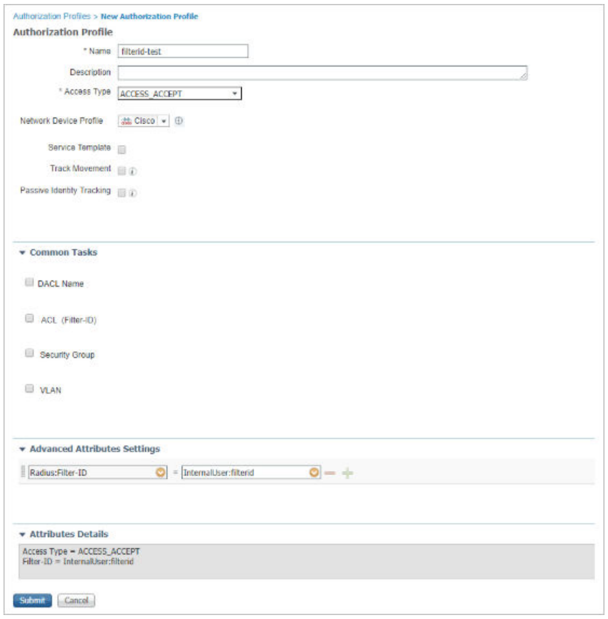

Authorization Profiles > New Authorization Profile

Authorization Profile

* Name: filterid-test

Description:

* Access Type: ACCESS_ACCEPT

Network Device Profile: Cisco

Service Template □

Track Movement □

Passive Identity Tracking □

Common Tasks

□ DACL Name

□ ACL (Filter-ID)

□ Security Group

□ VLAN

Advanced Attributes Settings

Radius:Filter-ID = InternalUser:filterid

Attributes Details

```
Access Type = ACCESS_ACCEPT
Filter-ID = InternalUser:filterid
```

For help with the parameters, see your Cisco ISE documentation, such as Authorization Profile Window.

Set Up with Juniper Mist Access Assurance

NOTE: As an Juniper Mist Access Assurance admin, use this procedure for general guidance only. For detailed and up-to-date information, see the Juniper Mist Access Assurance documentation and support site.

Navigate to **Organization** > **Auth Policy Labels** > **Add Label**.

Enter the parameters to create a role. Creating a role in Access Assurance is equivalent to a filter-id.

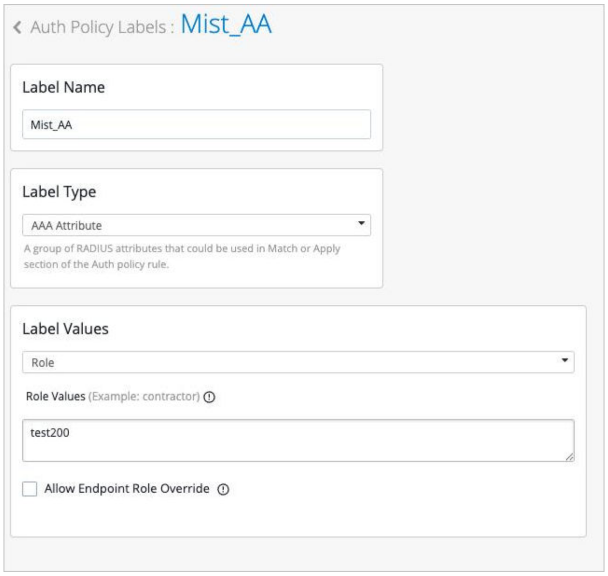

‹ Auth Policy Labels : Mist_AA

Label Name

Mist_AA

Label Type

AAA Attribute

A group of RADIUS attributes that could be used in Match or Apply section of the Auth policy rule.

Label Values

Role

Role Values (Example: contractor) ⓘ

test200

□ Allow Endpoint Role Override ⓘ

For help with the parameters, see Configure Authentication Policy Labels in the Juniper Mist Access Assurance Guide.

Assign a Profile to Unassigned Ports

You can specify a port profile that will be automatically assigned to all switch ports that are not associated with any specific port profile. Available at both the switch level and the template level, this feature provides greater flexibility and operational efficiency by ensuring consistent configuration of unassigned switch ports.

To set up automatic assignment of port profiles to unassigned ports:

1. Click **Organization** > **Wired** > **Switch Templates** to configure this feature at the organization level.

   Or, click **Site** > **Switch Configuration** to configure this feature at the site level.
2. Navigate to the **Select Switches Configuration** section.
3. Select a switch rule. Or, to add a new rule, click **Add Rule**.
4. Navigate to **Port Config** tab.

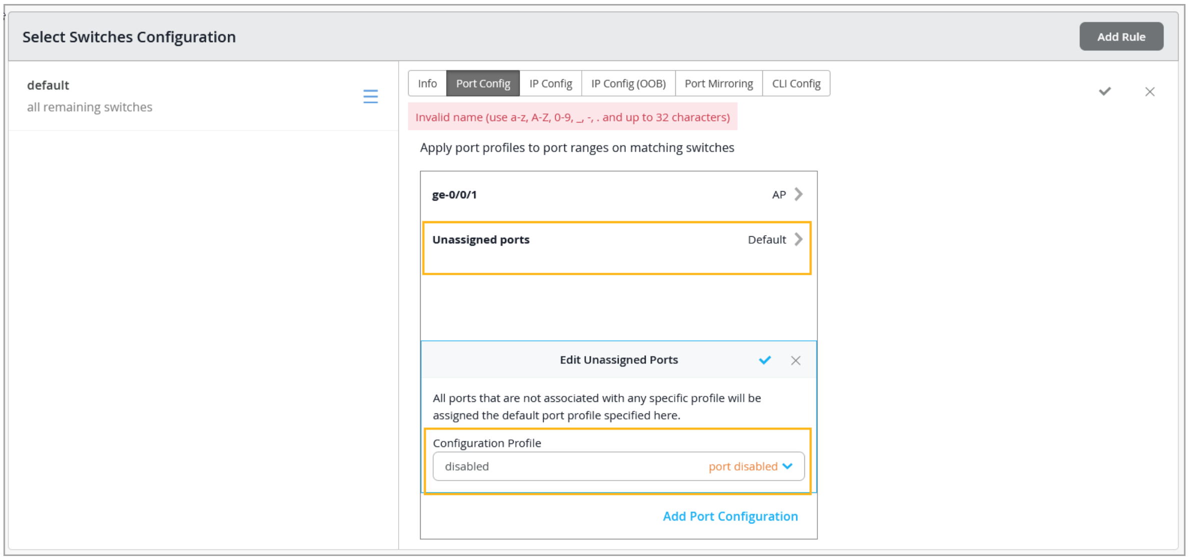

Invalid name (use a-z, A-Z, 0-9, _, -, . and up to 32 characters)

Apply port profiles to port ranges on matching switches

ge-0/0/1 AP

Unassigned ports Default

All ports that are not associated with any specific profile will be assigned the default port profile specified here.

Configuration Profile: disabled port disabled

5. Click the **Unassigned ports** tab.
6. From the **Configuration Profile** field, select a port profile and then click the check mark (✓) on the upper right of the **Edit Unassigned Ports** window.

   If you want to create a new profile and assign that to unassigned ports, use the **Add Port Configuration** option.
7. Click **Save** on the upper right of the page to save the configuration.

> **NOTE**: You can configure this feature at the switch level from the **Port** tile on the switch details page (**Switches** > *Switch Name*).

Override Port Configuration at the Switch Level

As an admin user, you can override a specific set of switch port configuration parameters at the switch level, giving you greater control over individual port settings. These overrides are appended to the existing port configuration and can be applied to one or more switch ports.

> **NOTE**: Port overrides are not permitted for ports that use dynamic port configuration.

To apply the override parameters at the switch level:

1. Click **Switches** and then select the switch to open the switch details page.
2. Navigate to **Port** tile.
3. On the Port tile, expand the **Port Overrides** section.

The Port Overrides section is visible only if at least one port profile configuration exists.

4. Click **Override Port Configuration**.

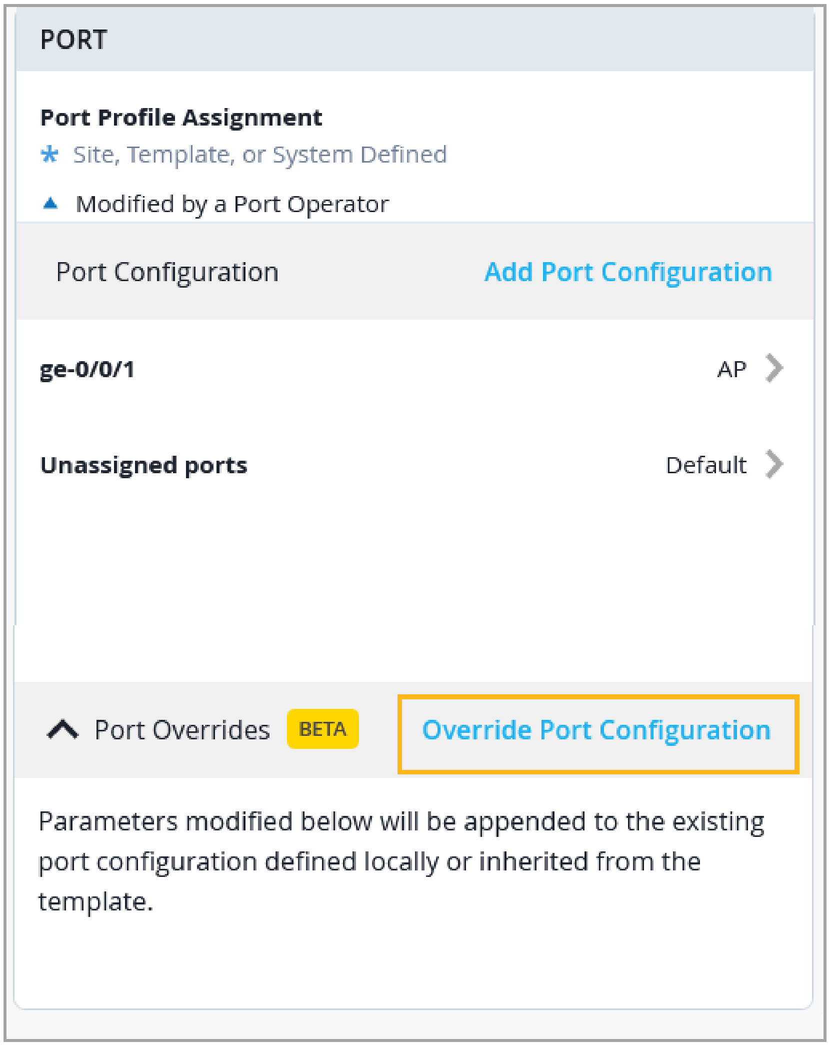

PORT

Port Profile Assignment

✱ Site, Template, or System Defined
▲ Modified by a Port Operator
Port Configuration
Add Port Configuration
ge-0/0/1
AP
Unassigned ports
Default
Port Overrides
BETA
Override Port Configuration

Parameters modified below will be appended to the existing port configuration defined locally or inherited from the template.

5. Specify the port IDs for which you want to configure the override parameters.
6. Select the parameters that you wish to override by checking the corresponding boxes.
7. Configure the selected parameters and then click the check mark (✓) on the right of the Add Port Overrides window.

   You can override the following parameters: Port Status, PoE, Speed, Duplex, MAC Limit, Port Description.

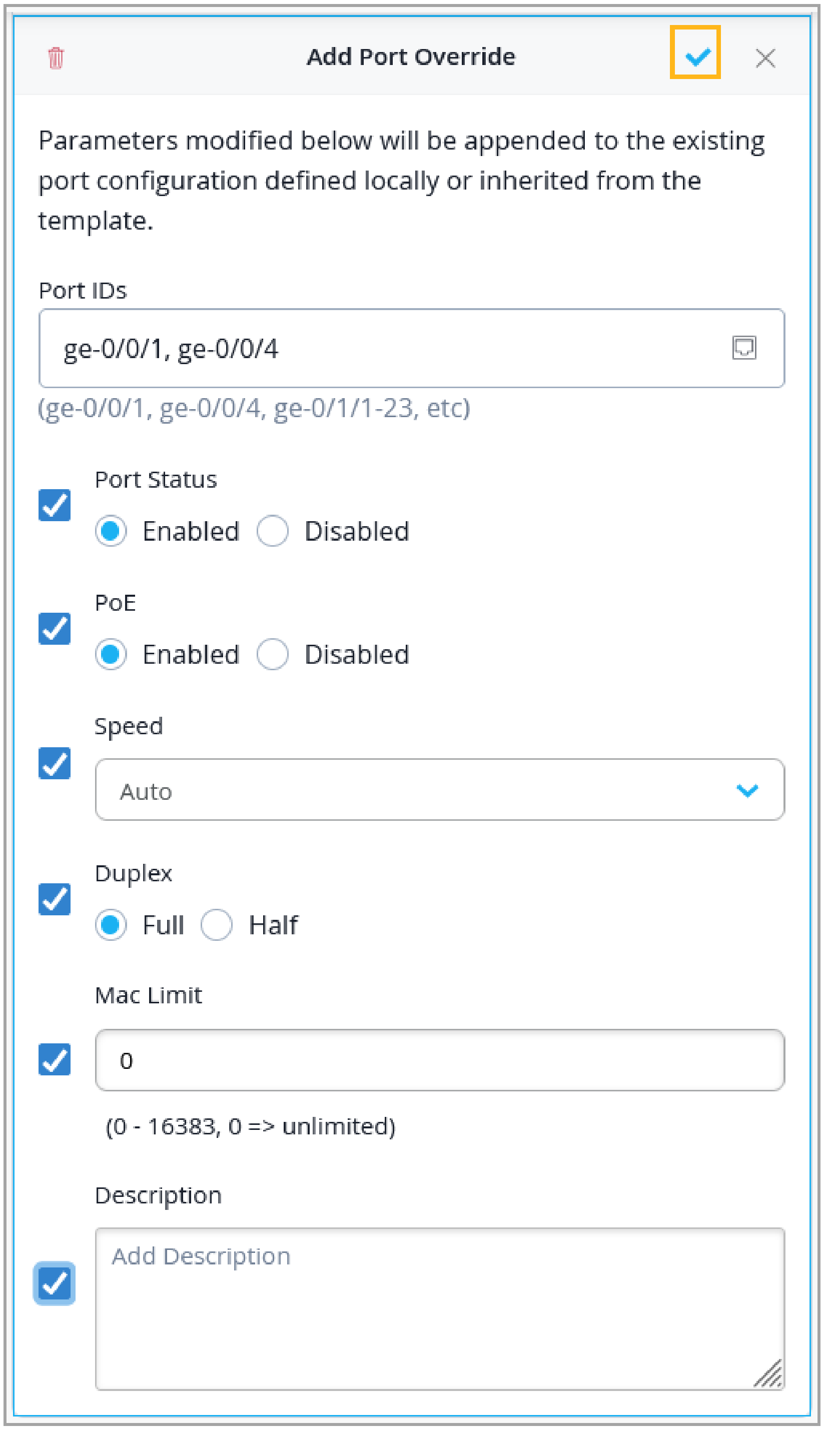

Add Port Override

Parameters modified below will be appended to the existing port configuration defined locally or inherited from the template.

Port IDs

ge-0/0/1, ge-0/0/4

(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

☑ Port Status

● Enabled ○ Disabled

☑ PoE

● Enabled ○ Disabled

☑ Speed

Auto

☑ Duplex

● Full ○ Half

☑ Mac Limit

0

(0 - 16383, 0 => unlimited)

☑ Description

Add Description

8. Click **Save** on the upper right of the switch details page to save the configuration.

Once applied, port overrides can be viewed in the Switch Ports section on the switch Insights page.

3
CHAPTER

Virtual Chassis Configuration

IN THIS CHAPTER

Virtual Chassis Overview (Juniper Mist)

SUMMARY

Learn about the benefits of using a Virtual Chassis, the VC support on various switch models, and the design guidelines for deploying VC at your site.

IN THIS SECTION

The Virtual Chassis technology enables you to connect multiple individual switches together to form one logical unit and to manage the individual switches as a single unit. You can configure and manage a Virtual Chassis using the Juniper Mist™ portal. The switches you add to a Virtual Chassis are called *members*. In a Virtual Chassis setup, Virtual Chassis ports (VCPs) connect the member switches and are responsible for passing the data and control traffic between member switches.

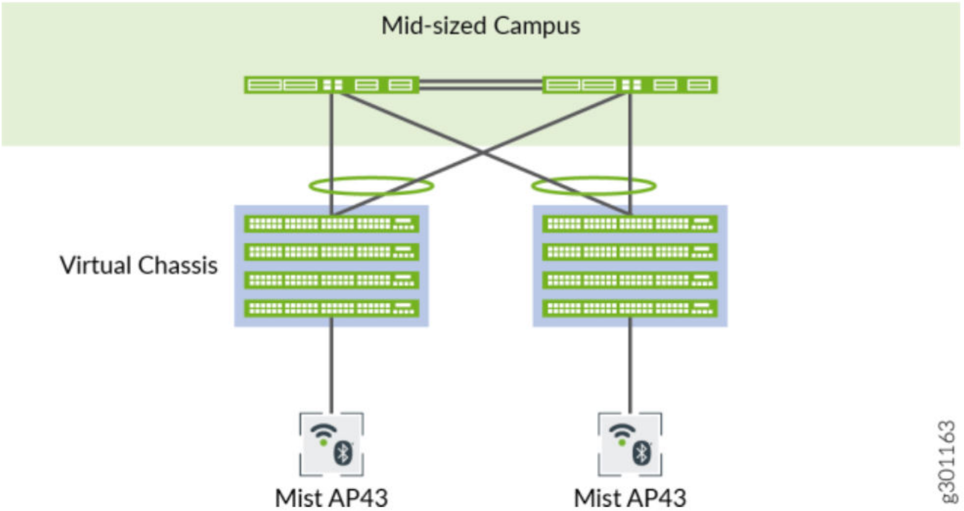

A Virtual Chassis helps you mitigate the risk of loops. It also eliminates the need for legacy redundancy protocols such as spanning tree protocols (STPs) and Virtual Router Redundancy Protocol (VRRP). In core and distribution deployments, you can connect to the Virtual Chassis using link aggregation group (LAG) uplinks. These uplinks ensure that the member switches in a Virtual Chassis have device-level redundancy.

A Virtual Chassis can include from two to 10 switches. Such a physical configuration can provide better resilience if a member switch goes down. One possible disadvantage to combining several switches into a Virtual Chassis is that this configuration requires more space and power than a single device requires.

Video: Virtual Chassis Overview

You can create a Virtual Chassis using the Form Virtual Chassis option on the Juniper Mist portal. The Form Virtual Chassis option applies only to the EX2300, EX4650, and QFX5120 switches as these switches don't have dedicated Virtual Chassis ports (VCPs). This option is not available to the EX3400, EX4100, EX4100-F, EX4300, and EX4400 switches as they come with dedicated VCPs. To create a Virtual Chassis with these switches, follow the procedure in this topic: "Configure a Virtual Chassis Using EX3400, EX4100, EX4100-F, EX4100-H, EX4300, EX4000, or EX4400" on page 212. However, the Modify Virtual Chassis workflow (see "Manage a Virtual Chassis Using Mist" on page 223) supports all the switches that Juniper supports Virtual Chassis on.

The table below shows the switch models along with the maximum number of member switches allowed in a Virtual Chassis configuration.

Table 18: Maximum Number of Member Switches Allowed in a Virtual Chassis Configuration

Switch Model	Maximum Member Switches Allowed
EX2300	4
EX4650	4
EX3400	10
EX4000	6
EX4100	10
EX4100-F	10
EX4100-H	6
EX4300	10
EX4400	10
QFX5120-32C, QFX5120-48T, QFX5120-48Y	2
QFX5120-48YM	4

Mist supports only preprovisioned Virtual Chassis configuration. It doesn't support nonprovisioned configuration. The preprovisioned configuration lets deterministically control the roles and member IDs assigned to the member switches when creating and managing a Virtual Chassis. The preprovisioned configuration distinguishes member switches by associating their serial numbers with the member ID.

For more information, see Virtual Chassis Overview for Switches

Mixed and Non-Mixed Virtual Chassis

A Virtual Chassis that includes switches of the same model can operate as a non-mixed Virtual Chassis. However, a Virtual Chassis that includes different models of the same switch (for example, two or more types of EX Series switches) must operate in mixed mode because of architecture differences between the different switch models.

Table 19: Supported Combination of Switches in a Mixed-Mode Virtual Chassis

Allowed Routing Engine Members	Allowed Linecard Members
EX4300	EX4300 and EX4600
EX4300-48MP	EX4300-48MP and EX4300 (excludes EX4600)
EX4600	EX4600 and EX4300 (excludes EX4300-48MP)

ℹ **NOTE**: Currently, Mist cloud does not support creation of a Virtual Chassis using the EX4600 switches.

For more information about the combination of switches that a mixed or a non-mixed Virtual Chassis configuration supports, see Understanding Mixed EX Series and QFX Series Virtual Chassis.

Design Considerations for Virtual Chassis

We recommend that you physically distribute your Juniper access points (APs) across the network operations center (NOC) floor. This helps you connect each switch in a virtual stack to a different AP. Doing so provides better redundancy. This design also helps in handling hardware failures related to power supply.

Figure 7: Virtual Chassis Setup in a NOC

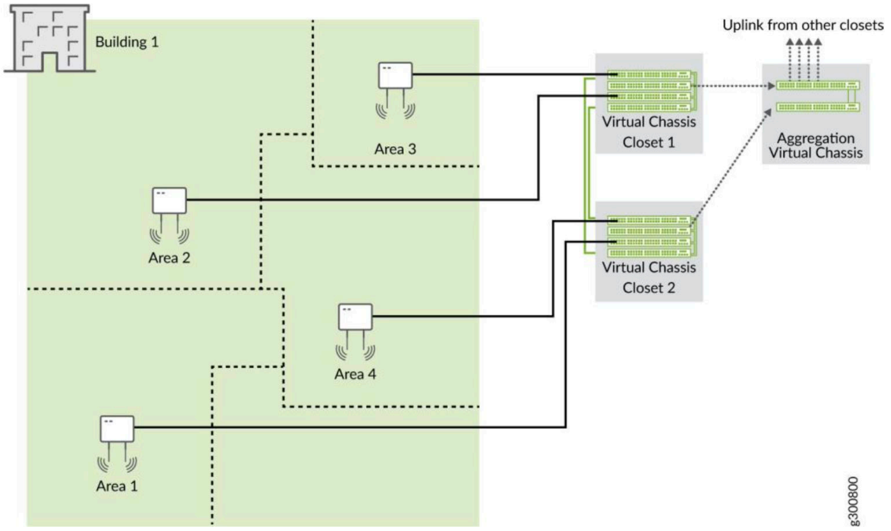

For example, you can use one of the following two options if you want to deploy a solution that includes 96 ports:

- Use two EX4300-48P switches, with one switch serving as the primary and the other as the backup. This option is cost effective and ensures a compact footprint. The main disadvantage of this option is that a failure of one switch can negatively affect 50 percent of your users.

- Use four EX4300-24P switches, with one switch serving as the primary, one as the backup, and the remaining two switches as linecard members. This option provides a better high availability because any failure of one switch affects only 25 percent of users. A switch failure does not necessarily affect the uplinks (if the failed switch did not include any uplinks). This option requires more space and power.

Regardless of the options you choose, we recommend that you do the following:

- Configure the primary and backup switches in the Virtual Chassis in such a way that they are in different physical locations.

- Distribute the member switches of the Virtual Chassis in such a way that no more than half of the switches depend on the same power supply or any single point of failure.

- Space the member switches evenly by a member hop in the Virtual Chassis.

Configure a Virtual Chassis Using EX3400, EX4100, EX4100-F, EX4100-H, EX4300, EX4000, or EX4400

SUMMARY

Follow these steps to configure a Virtual Chassis on Juniper switches with a dedicated Virtual Chassis port.

NOTE:

- Instructions for configuring a Virtual Chassis using the EX4400-24X switches are slightly different. For details, refer to "Configure a Virtual Chassis Using EX4400-24X" on page 221.
- Ensure that you use the Mist portal interfaces to manage all Virtual Chassis configurations. Avoid using CLIs or Additional CLIs for managing Virtual Chassis settings.
- Juniper Mist automatically upgrades a Virtual Chassis linecard member if it is running a Junos version different from that of the primary member. The linecard member will be upgraded to the same version as the primary member if the following conditions are met:
    - The switch must form a Virtual Chassis with three or more members—that is, a primary, a backup, and a linecard member.
    - The Junos version on the linecard member is different from that on the primary member.
    - The linecard member must be in Inactive state.

      Note that a linecard member will be upgraded only if it is inactive and running a clearly different Junos version. Minor differences, such as different spin numbers, will not trigger an upgrade.

- Only the Junos versions listed on the Mist portal are available for upgrade.

The EX3400, EX4100, EX4100-F, EX4100-H, EX4300, EX4000, and EX4400 switches come with dedicated Virtual Chassis ports (VCPs). To create Virtual Chassis using these switches, you only need to connect them to each other via VCPs. The **Form Virtual Chassis** option on the Switches page on the Mist portal is not applicable to these switches. However, once a Virtual Chassis is created with these switches, you can use the Modify Virtual Chassis option on the switch details page to modify and manage the Virtual Chassis. The Virtual Chassis workflow for these switches involves the following two steps:

- Virtual Chassis formation by connecting the switches via the dedicated VCPs and powering on them.
- Preprovisioning the Virtual Chassis using the Modify Virtual Chassis option on the Juniper Mist Portal. Mist supports only the preprovisioned Virtual Chassis configuration. The preprovisioned configuration specifies the chassis serial number, member ID, and role for both member switches in the Virtual Chassis. When a new member router joins the Virtual Chassis, Junos compares its serial number against the values specified in the preprovisioned configuration. Preprovisioning prevents any accidental role assignments, or the accidental addition of a new member to the Virtual Chassis. Each role, member ID, addition or removal of members, is under the control of the configuration.

In addition to Virtual Chassis creation, you can renumber, replace, or add a member to an existing Virtual Chassis, by using the Modify Virtual Chassis option on the Switch Details Page.

You can configure the Virtual Chassis in mixed mode or non-mixed mode. A Virtual Chassis that includes switches of the same model operates as a non-mixed Virtual Chassis. However, a Virtual Chassis that includes different models of the same switch operates in mixed mode because of architecture differences between the different switch models. For more information, see "Mixed and Non-Mixed Virtual Chassis" on page 210.

Table 20: Supported Combination of Switches in a Mixed-Mode Virtual Chassis

Allowed Routing Engine Members	Allowed Linecard Members
EX4300	EX4300
EX4300-48MP	EX4300-48MP and EX4300

To configure a Virtual Chassis using EX3400, EX4100, EX4100-F, EX4300, or EX4400 switches:

1. Ensure that all the switches that you want to include in the Virtual Chassis are onboarded to the Juniper Mist cloud and assigned to the same site. To onboard a new switch (greenfield deployment), see Onboard Switches to Mist Cloud. To onboard an existing switch (brownfield deployment), see "Onboard a Brownfield Switch" on page 18.
2. Power off the switches that you want to include in the Virtual Chassis.

3. Connect the switches to each other using the dedicated Virtual Chassis ports (VCPs), preferably in a full ring topology, as shown below. The following is a sample image. The location of the VCPs will vary depending on the switch models.

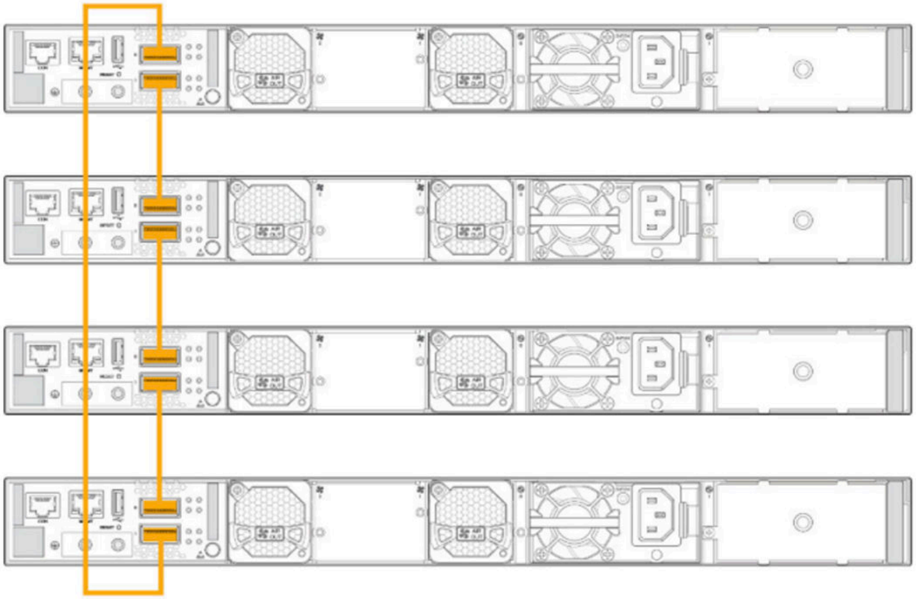

NOTE:
On EX4100-F-12P and EX4100-F-12T switches, the 10-Gigabit Ethernet ports in PIC 2, which are used as uplink ports, cannot be configured as Virtual Chassis ports. These ports do not support the HGoE mode. For more information about the HGoE mode, see HiGig and HGoE Modes.

4. Power on the switches.
5. Wait for the MST LED on the primary and backup switches to come up. The LED appears solid on the primary switch. On the backup switch, the LED stays in a blinking state.

NOTE: MST LED remains Off on the switches elected as the Linecard members in a Virtual Chassis.

A Virtual Chassis is now physically formed but not preprovisioned.
6. Connect the Virtual Chassis to the Juniper Mist cloud by connecting the uplink port on the primary switch to the upstream switch.

We recommend connecting the uplink port only after the Virtual Chassis has formed. Wait for the MST LEDs to come up (LED appears solid on the primary member and blinking on the backup member), then connect the uplink ports on those members.

This step initiates a zero-touch provisioning (ZTP) process on the Virtual Chassis and connects it to the Juniper Mist cloud.

After a Virtual Chassis connects to the Mist cloud for the first time, it may take 5 to 10 minutes for the Virtual Chassis stats to appear on the Mist cloud for all the members.

7. Click **Switches** > *Switch Name* to go to the Virtual Chassis page (the switch details page) to verify the details.

   The switches appear as a single Virtual Chassis as shown below:

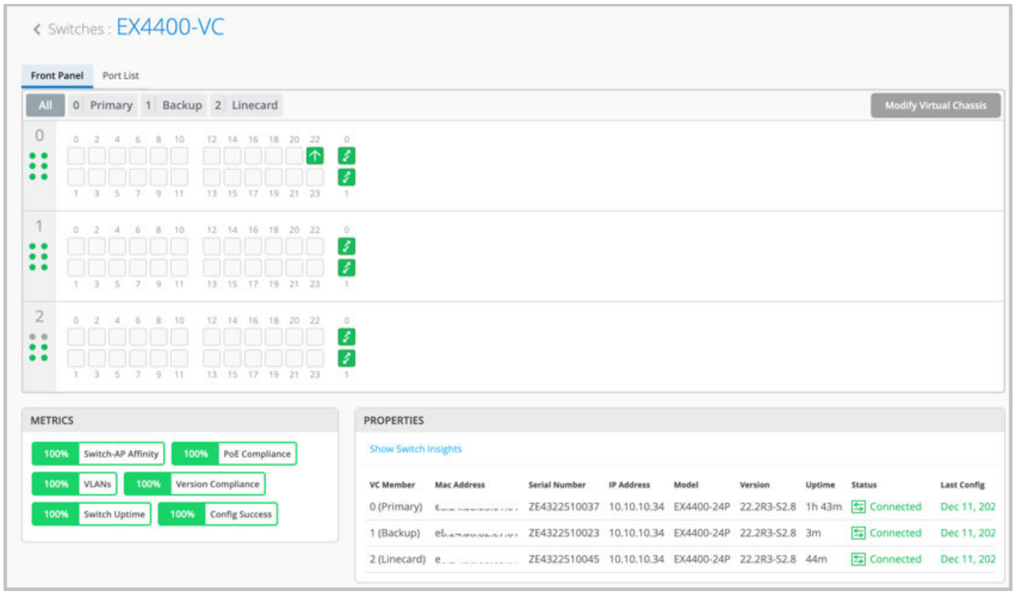

< Switches : EX4400-VC

METRICS

100%	Switch-AP Affinity
100%	PoE Compliance
100%	VLANs
100%	Version Compliance
100%	Switch Uptime
100%	Config Success

PROPERTIES

Show Switch Insights

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last Config
0 (Primary)	[redacted]	ZE4322510037	10.10.10.34	EX4400-24P	22.2R3-S2.8	1h 43m	Connected	Dec 11, 202
1 (Backup)	[redacted]	ZE4322510023	10.10.10.34	EX4400-24P	22.2R3-S2.8	3m	Connected	Dec 11, 202
2 (Linecard)	[redacted]	ZE4322510045	10.10.10.34	EX4400-24P	22.2R3-S2.8	44m	Connected	Dec 11, 202

8. After Virtual Chassis is connected to the Mist cloud, preprovision it. Preprovisioning allows users to define the roles and renumber appropriately. To preprovision the Virtual Chassis, follow the steps below:

   a. On the switch details page, click **Modify Virtual Chassis**.

      The Modify Virtual Chassis page appears.

      > **NOTE**: The Modify Virtual Chassis option is available for switches that have the configuration management enabled in Mist.

   b. On the Modify Virtual Chassis page, click **Preprovision Virtual Chassis**.

      This step pushes the preprovisioned Virtual Chassis configuration to the device and overwrites the old autoprovision Virtual Chassis configuration pushed to the device during the ZTP process. This option assumes the current positioning of the members and preprovisions them as is.

> **NOTE**: If you make any changes on the Modify Virtual Chassis page, such as moving the members around or adding or removing members, the Preprovision Virtual Chassis

button is disabled and the Update button is enabled. In this case, click the **Update** button to effect the changes made and Preprovision the Virtual Chassis.

All configurations are pushed instantly after you preprovision the Virtual Chassis. The stats could take up to 15 minutes to appear on the Mist dashboard.

For a detailed procedure on how to modify a Virtual Chassis, see "Manage a Virtual Chassis Using Mist" on page 223.

Configure a Virtual Chassis Using EX2300, EX4650, or QFX5120 Switches

SUMMARY

Follow these steps to configure Virtual Chassis on Juniper switches that don't have a dedicated VC port. Note that the screenshots used in this procedure are for reference only.

The Juniper Networks EX2300, EX4650, and QFX5120 switches do not form a Virtual Chassis by default, as these switches don’t have dedicated Virtual Chassis ports (VCPs). Therefore, to create a Virtual Chassis with these switches, you need to use the Form Virtual Chassis option on the Juniper Mist™ portal. The Form Virtual Chassis option applies only to the EX2300, EX4650, and QFX5120 switches. This workflow creates a preprovisioned Virtual Chassis configuration. Mist supports only the preprovisioned Virtual Chassis configuration.

The procedure to configure a Virtual Chassis using the EX3400, EX4100, EX4100-F, EX4300, or EX4400 switches is different, as those switches have dedicated Virtual Chassis ports (VCPs). For more information, see "Configure a Virtual Chassis Using EX3400, EX4100, EX4100-F, EX4100-H, EX4300, EX4000, or EX4400" on page 212.

NOTE: Ensure that you use the Mist portal interfaces to manage all Virtual Chassis configurations. Avoid using CLIs or Additional CLIs for managing Virtual Chassis settings.

To configure a Virtual Chassis using EX2300, EX4650, or QFX5120 switches:

1. Connect the switches to the Mist cloud. Ensure that you have an uplink connection directly to the switch.

2. Ensure that all the switches that you want to include in the Virtual Chassis are onboarded to the Juniper Mist™ cloud and assigned to the same site. Also, ensure that configuration management is enabled on the switches.

   To onboard a new switch (greenfield deployment), see Onboard Switches to Mist Cloud. To onboard an existing switch (brownfield deployment) to Juniper Mist, see "Onboard a Brownfield Switch" on page 18.

3. Click the **Switches** tab on the left to navigate to the Switches page.

4. Select the switches that you want to include in the Virtual Chassis.

   An EX2300 switch variant can form a Virtual Chassis with any EX2300 switch variants. An EX4650 variant switch can form a Virtual Chassis with any EX4650 switch variants. A QFX5120 switch variant can form a Virtual Chassis only with the same QFX5120 switch variant. Therefore, the **Form Virtual Chassis** option is available only if you select the right switch models for a Virtual Chassis.

5. Click **More** > **Form Virtual Chassis**.

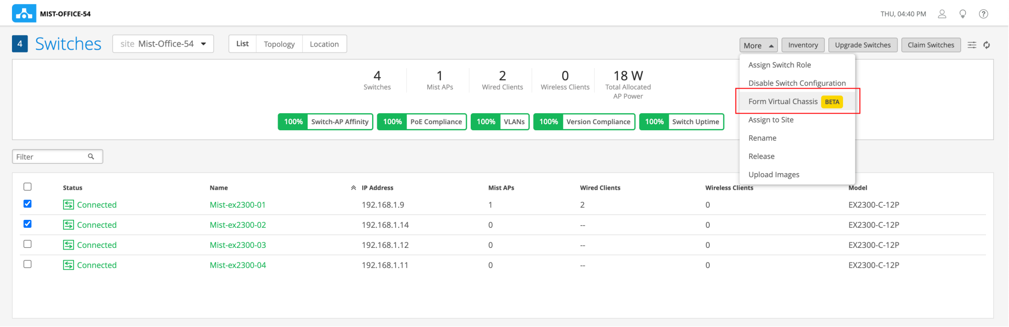

□	Status	Name	IP Address	Mist APs	Wired Clients	Wireless Clients	Model
☑	Connected	Mist-ex2300-01	192.168.1.9	1	2	0	EX2300-C-12P
☑	Connected	Mist-ex2300-02	192.168.1.14	0	--	0	EX2300-C-12P
□	Connected	Mist-ex2300-03	192.168.1.12	0	--	0	EX2300-C-12P
□	Connected	Mist-ex2300-04	192.168.1.11	0	--	0	EX2300-C-12P

NOTE: You can see the **Form Virtual Chassis** option only if:

- The selected switches are running the same Junos version and have the configuration management option enabled.
- All the selected switch models are supported by the Virtual Chassis.

You can also create Virtual Chassis from the switch details page by using the **Utilities** > **Form Virtual Chassis** option.

The Form Virtual Chassis window appears, as shown in the following example.

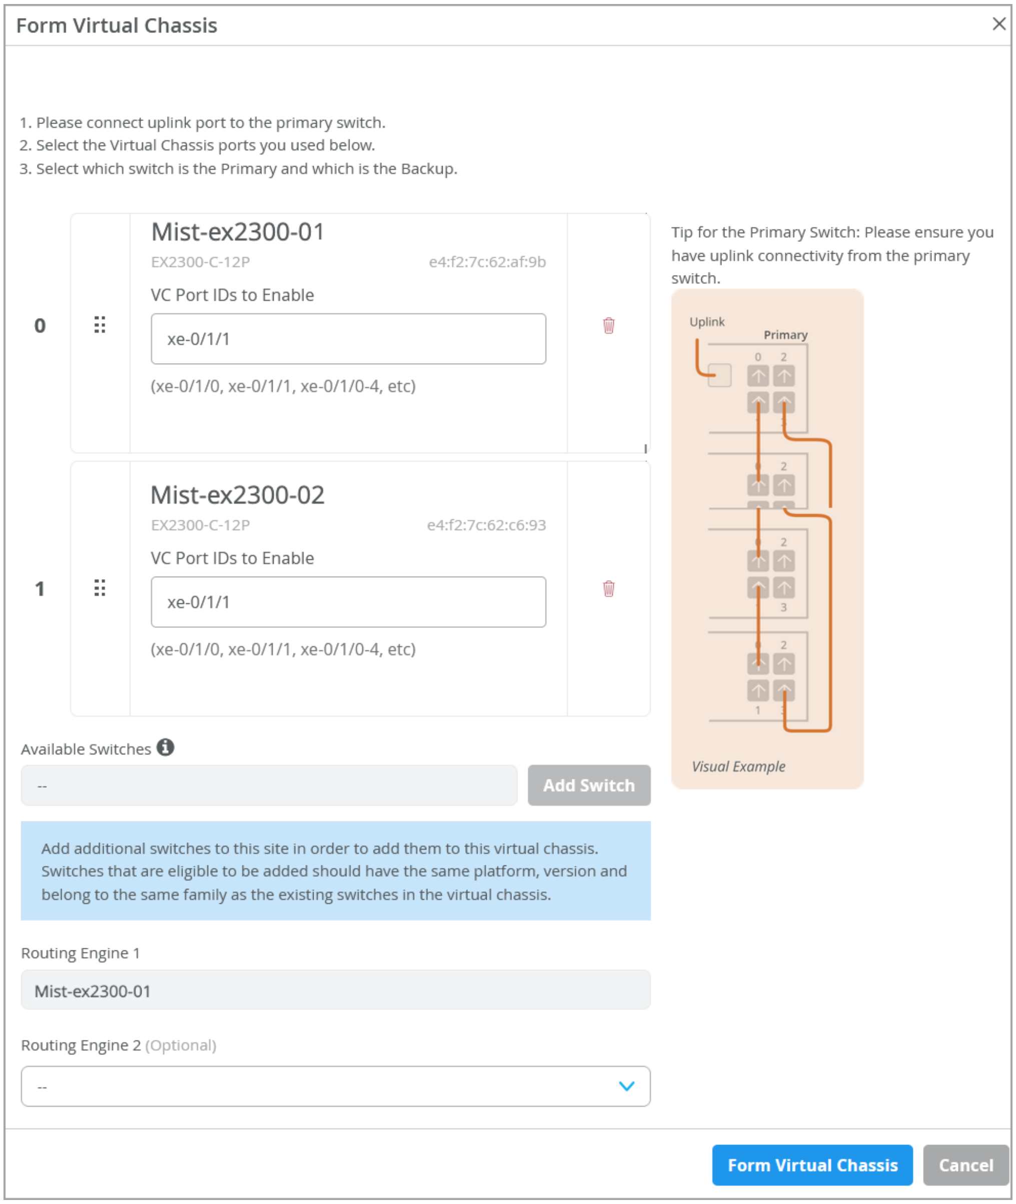

Form Virtual Chassis

1. Please connect uplink port to the primary switch.
2. Select the Virtual Chassis ports you used below.
3. Select which switch is the Primary and which is the Backup.

0 Mist-ex2300-01
EX2300-C-12P e4:f2:7c:62:af:9b
VC Port IDs to Enable: xe-0/1/1
(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

1 Mist-ex2300-02
EX2300-C-12P e4:f2:7c:62:c6:93
VC Port IDs to Enable: xe-0/1/1
(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

Tip for the Primary Switch: Please ensure you have uplink connectivity from the primary switch.

Available Switches: --

Add additional switches to this site in order to add them to this virtual chassis. Switches that are eligible to be added should have the same platform, version and belong to the same family as the existing switches in the virtual chassis.

Routing Engine 1: Mist-ex2300-01

Routing Engine 2 (Optional): --

> **NOTE**: This example shows two switches included in the Virtual Chassis. A Virtual Chassis device created using EX2300 or EX4650 switches supports up to 4 switches. All switches, except those assigned Routing Engine roles, function as linecard members.

6. On the Form Virtual Chassis window, specify the following:

   a. **Port IDs** for the switches. These are IDs for the Virtual Chassis ports (VCPs). This window displays all the switches you selected from the Switches page.

   b. The **Routing Engine 1** switch. The switch that you selected first is shown as the Routing Engine 1 switch by default. You can modify that.

c. The **Routing Engine 2** switch. This configuration is optional. If you don't select a switch to function in the Routing Engine 2 role, Mist assigns the linecard role to that switch.

NOTE:

- Juniper Mist automatically upgrades a Virtual Chassis linecard member if it is running a Junos version that is different from that of the primary member. The linecard member will be upgraded to the same version as the primary member if the following conditions are met:
  - The switch must form a Virtual Chassis with three or more members—that is, a primary, a backup, and a linecard member.
  - The Junos version on the linecard member is different from that on the primary member.
  - The linecard member must be in Inactive state.

    Note that a linecard member will be upgraded only if it is inactive and running a clearly different Junos version. Minor differences, such as different spin numbers, will not trigger an upgrade.
  - Only the Junos versions listed on the Mist portal are available for upgrade.
- Ensure that port configurations are tailored to the number of members in the Virtual Chassis. For example, the following configuration should be applied only to a 5-member Virtual Chassis:

  ```
  ge-0/0/0-47,
  ge-0/0/1-47
  ge-0/0/2-47
  ge-0/0/3-47
  ge-0/0/4-47
  ```

  Avoid applying the same port settings across Virtual Chassis devices of different sizes.

This operation converts the ports to Virtual Chassis ports (VCPs) and also pre-provisions the Virtual Chassis.

7. Click **Form Virtual Chassis** and wait for the Virtual Chassis to be created.

   The switches page shows a message indicating that you must connect the switches to each other using the VCPs.

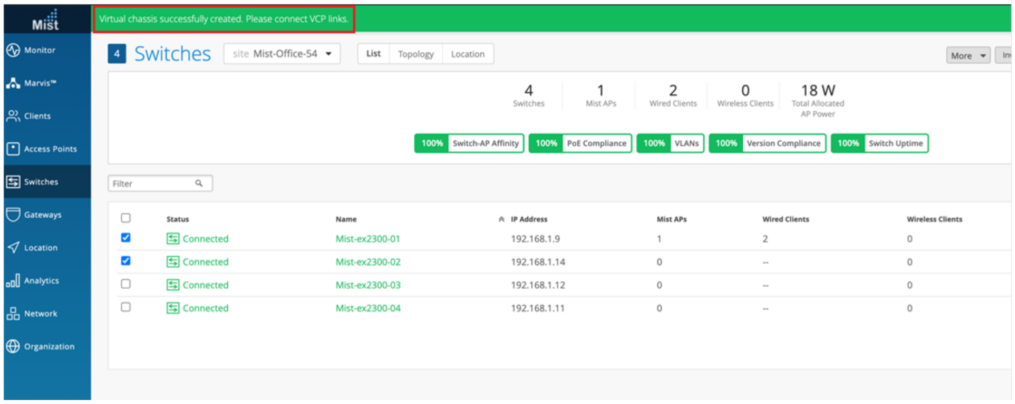

Virtual chassis successfully created. Please connect VCP links.

☐	Status	Name	IP Address	Mist APs	Wired Clients	Wireless Clients
☑	Connected	Mist-ex2300-01	192.168.1.9	1	2	0
☑	Connected	Mist-ex2300-02	192.168.1.14	0	--	0
☐	Connected	Mist-ex2300-03	192.168.1.12	0	--	0
☐	Connected	Mist-ex2300-04	192.168.1.11	0	--	0

8. Connect the switches to each other using the VCPs mentioned earlier.

   When the Virtual Chassis formation is in progress, the Switches page shows the switch status as **VC forming**.

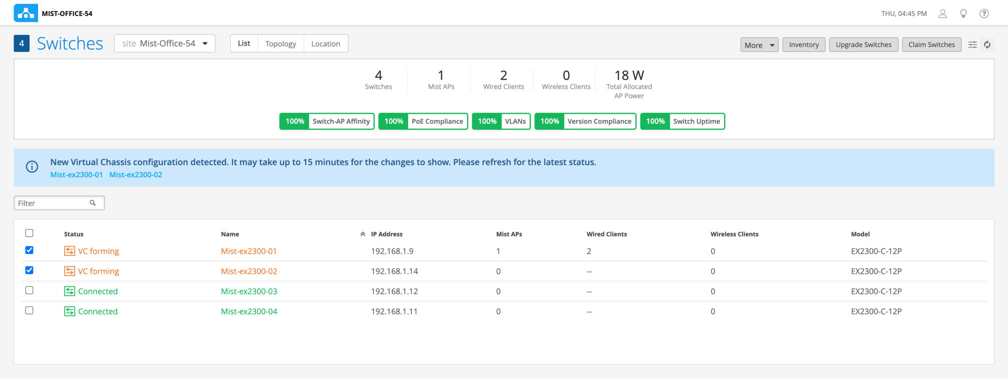

New Virtual Chassis configuration detected. It may take up to 15 minutes for the changes to show. Please refresh for the latest status.
Mist-ex2300-01 Mist-ex2300-02

☐	Status	Name	IP Address	Mist APs	Wired Clients	Wireless Clients	Model
☑	VC forming	Mist-ex2300-01	192.168.1.9	1	2	0	EX2300-C-12P
☑	VC forming	Mist-ex2300-02	192.168.1.14	0	--	0	EX2300-C-12P
☐	Connected	Mist-ex2300-03	192.168.1.12	0	--	0	EX2300-C-12P
☐	Connected	Mist-ex2300-04	192.168.1.11	0	--	0	EX2300-C-12P

After the Virtual Chassis formation is successful, the Switches page displays only one entry for the Virtual Chassis with the name of the primary switch. However, the MIST APs column displays one AP for each Virtual Chassis member in a comma-separated format.

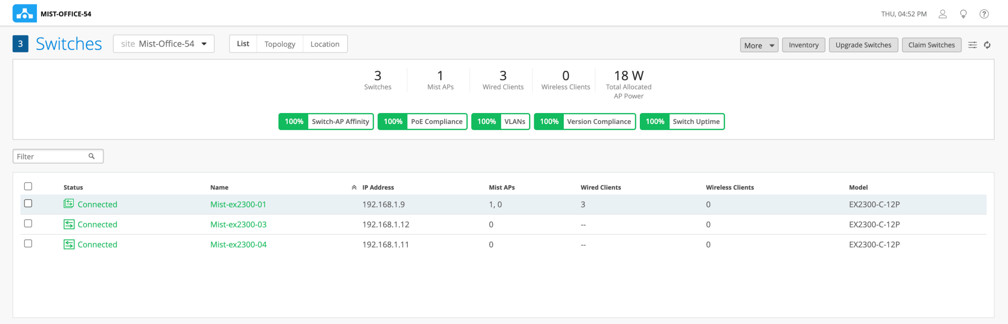

☐	Status	Name	IP Address	Mist APs	Wired Clients	Wireless Clients	Model
☐	Connected	Mist-ex2300-01	192.168.1.9	1, 0	3	0	EX2300-C-12P
☐	Connected	Mist-ex2300-03	192.168.1.12	0	--	0	EX2300-C-12P
☐	Connected	Mist-ex2300-04	192.168.1.11	0	--	0	EX2300-C-12P

The switch details page displays the front panel of all the Virtual Chassis members.

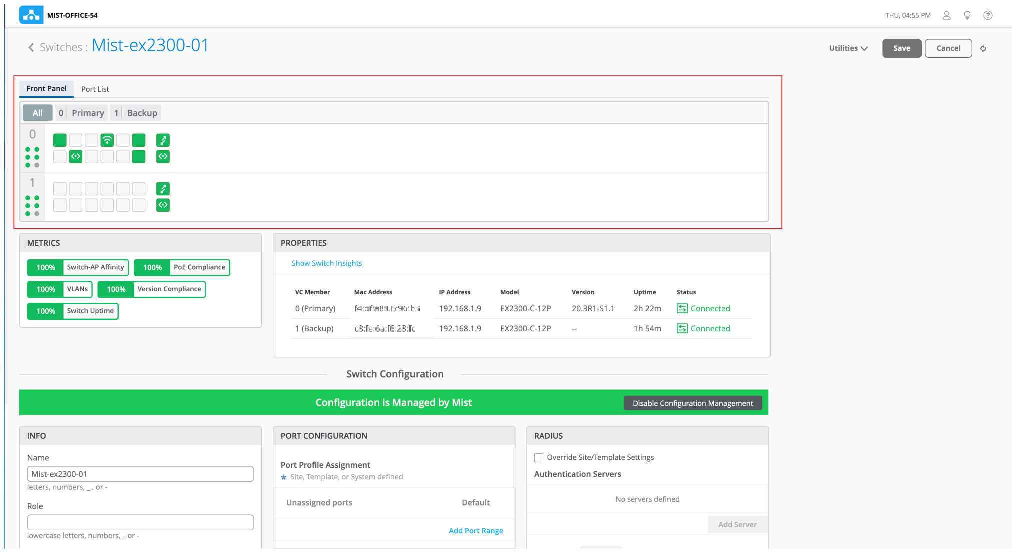

VC Member	Mac Address	IP Address	Model	Version	Uptime	Status
0 (Primary)	[redacted]	192.168.1.9	EX2300-C-12P	20.3R1-S1.1	2h 22m	Connected
1 (Backup)	[redacted]	192.168.1.9	EX2300-C-12P	--	1h 54m	Connected

Name: Mist-ex2300-01
letters, numbers, _ . or -

Role:
lowercase letters, numbers, _ or -

Port Profile Assignment
* Site, Template, or System defined

Unassigned ports: Default

☐ Override Site/Template Settings

Authentication Servers

No servers defined

NOTE: Once the Virtual Chassis is formed, if you need only one uplink to the Virtual Chassis, maintain the uplink to primary switch and remove uplinks from other switches.

You can use the **Modify Virtual Chassis** option on the switch details page to renumber and replace Virtual Chassis members and add members to a Virtual Chassis connected to the Mist cloud. For more information, see "Manage a Virtual Chassis Using Mist" on page 223.

Configure a Virtual Chassis Using EX4400-24X

NOTE:

- Virtual Chassis formation on the EX4400-24X switches cannot be executed without pre-provisioning. You may need to enable this in a lab or staging environment where you have proper access to the Mist cloud managing the devices so that they appear on the **Inventory** page. For EX4400-24X switches, plan to have serial console access when pre-provisioning the Virtual Chassis.

- Juniper Mist automatically upgrades a Virtual Chassis linecard member if it is running a Junos version different from that on the primary member. The linecard member will be upgraded to the same version as that on the primary member if the following conditions are met:

- The switch must form a Virtual Chassis with three or more members—that is, a primary, a backup, and a linecard member.
- The Junos version on the linecard member is different from that on the primary member.
- The linecard member must be in Inactive state.

  Note that a linecard member will be upgraded only if it is inactive and running a clearly different Junos version. Minor differences, such as different spin numbers, will not trigger an upgrade.
- Only the Junos versions listed on the Mist portal are available for upgrade.

- Ensure that you use the Mist portal interfaces to manage all Virtual Chassis configurations. Avoid using CLIs or Additional CLIs for managing Virtual Chassis settings.

To configure a Virtual Chassis using EX4400-24X follow the steps below (in these steps, we use two EX4400-24X switches as an example):

1. Unbox and power on the EX4400-24X switches.
2. For cloud reachability, connect the 10GbE front panel ports or insert an appropriate uplink module on each EX4400-24X Virtual Chassis member to the upstream devices, such as the WAN router. Ensure that the links are up and can reach the Juniper Mist cloud.
3. Onboard the switch to the Mist cloud, using the claim or adopt method. For more information, refer to Onboard Switches to Mist Cloud.

   When using the claim method, the devices should appear in the Inventory of the Juniper Mist cloud automatically as part of the process. Using the adopt method, the switches may appear immediately or after the Virtual Chassis has been formed.
4. Connect the Virtual Chassis ports (VCPs) on EX4400-24X using DAC Cables – 40G/100G. The Virtual Chassis ports are located on the front-panel on the EX4400-24X (as indicated in the below figure) and only support the HGoE protocol for Virtual Chassis formation.

   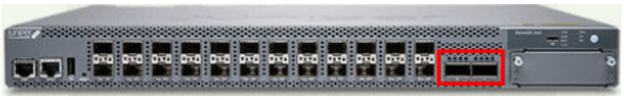
5. Using either Remote Console or a serial console cable connected directly to the first switch that will become the primary switch of your Virtual Chassis, log in to the Junos CLI by entering the following command:

```
request virtual-chassis mode hgoe <reboot>
```

This CLI converts the two front panel ports on the EX4400-24X switches from network ports to VCPs to enable Virtual Chassis formation. After entering this command, ensure that the switch is rebooted.

6. Repeat the above CLI command on your backup switch and then reboot it. Also, repeat these steps on any optional linecard switches.

7. Once the devices come up on the Mist portal after the reboot, they should be in Virtual Chassis mode and seen as a single device on the Mist dashboard. You can verify this from the switch details page by clicking **Switches** > ***Switch Name***.

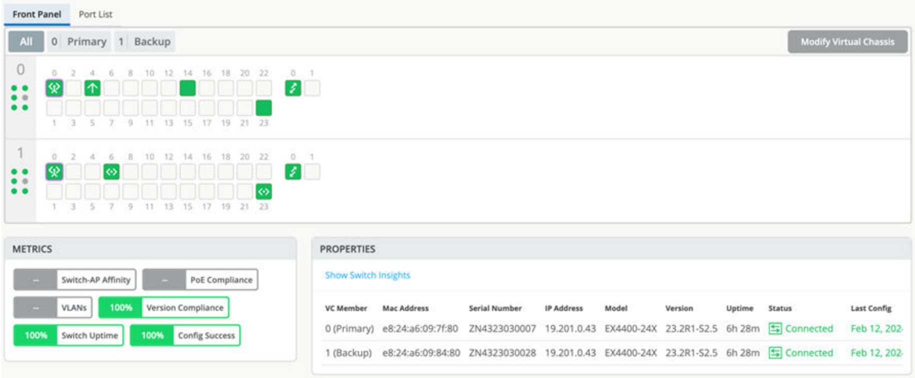

--	Switch-AP Affinity
--	PoE Compliance
--	VLANs
100%	Version Compliance
100%	Switch Uptime
100%	Config Success

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last Config
0 (Primary)	e8:24:a6:09:7f:80	ZN4323030007	19.201.0.43	EX4400-24X	23.2R1-S2.5	6h 28m	Connected	Feb 12, 202-
1 (Backup)	e8:24:a6:09:84:80	ZN4323030028	19.201.0.43	EX4400-24X	23.2R1-S2.5	6h 28m	Connected	Feb 12, 202-

8. Optional: Check the Virtual Chassis state via Remote Shell as shown in the figure below:

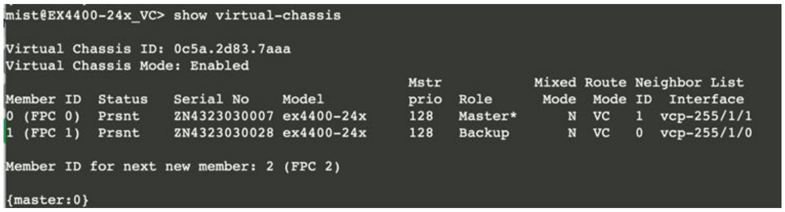

```
mist@EX4400-24x_VC> show virtual-chassis

Virtual Chassis ID: 0c5a.2d83.7aaa
Virtual Chassis Mode: Enabled
                                                Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model          prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    ZN4323030007 ex4400-24x     128   Master*      N  VC     1  vcp-255/1/1
1 (FPC 1)  Prsnt    ZN4323030028 ex4400-24x     128   Backup       N  VC     0  vcp-255/1/0

Member ID for next new member: 2 (FPC 2)

{master:0}
```

To set up a Virtual Chassis using any EX4400 switch models other than EX4400-24X, refer to "Configure a Virtual Chassis Using EX3400, EX4100, EX4100-F, EX4100-H, EX4300, EX4000, or EX4400" on page 212.

Manage a Virtual Chassis Using Mist

SUMMARY

Follow these steps to manage the members of a Virtual Chassis.

IN THIS SECTION

You can use the **Modify Virtual Chassis** option on the switch details page to manage your Virtual Chassis. The operations you can perform include renumbering and replacing the Virtual Chassis members and adding new members to a Virtual Chassis.

The Modify Virtual Chassis workflow leverages the Junos preprovisioning method which configures the role and serial number of all members in a Virtual Chassis.

NOTE: The Modify Virtual Chassis option is available for switches that have the configuration management enabled in Mist.

The preprovisioned configuration specifies the chassis serial number, member ID, and role for both member switches in the Virtual Chassis. When a new member router joins the Virtual Chassis, Junos compares its serial number against the values specified in the preprovisioned configuration. Preprovisioning prevents any accidental role assignment to a Routing Engine, or any accidental addition of a new member to the Virtual Chassis. Role assignments, member ID assignments, and additions or deletions of members in Virtual Chassis are under the control of a preprovisioned configuration.

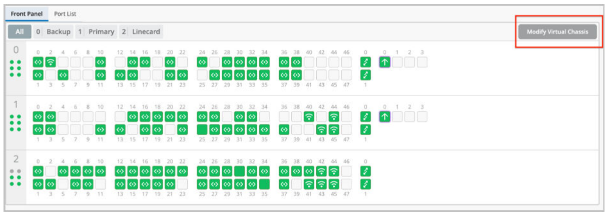

NOTE:

- The Modify Virtual Chassis option is available:
  - To Super Users or Network Admins.
  - For switches that have their configuration managed by Mist.
- This workflow applies to all the EX Series and QFX Series platforms that support Virtual Chassis.
- To delete a member whose MAC address is used as the Virtual Chassis device ID, trash and replace it with an existing member in the Virtual Chassis. To verify if any Virtual Chassis member is used as the device identifier, look for the device ID on the switch details page (Virtual Chassis page) or on the switch list.
- The Add Switch dropdown only shows the switches that:
  - Are part of the same site. Models with dedicated Virtual Chassis ports can be in connected or disconnected state. However, to modify the EX2300, EX4650, or QFX5120 Virtual Chassis, the members should be in the connected state as these switches don't have dedicated Virtual Chassis ports.
  - Have configuration management enabled in Mist.
  - Are not currently part of the same or another Virtual Chassis.
  - Are of the same model family. For example, an EX4100-F switch can be part of a Virtual Chassis with an EX4100-48MP switch.
- The Modify Virtual Chassis button is disabled when the Configuration Management option is disabled for the switch.
- When a Virtual Chassis configuration is in progress, you cannot make any changes inside the Modify Virtual Chassis page.

Prerequisites

Before your perform any modification to a Virtual Chassis, you must remove all the additional CLI commands specific to Virtual Chassis (the `virtual-chassis` commands) from the associated device or site template. The additional CLI commands take precedence over other types of configurations. If a Virtual Chassis configuration is detected under the Additional CLI Commands section, you cannot make any changes using the **Modify Virtual Chassis** option. When you attempt to modify a Virtual Chassis, the

Mist dashboard displays a message to indicate that the Additional CLI commands (if present) need to be removed and saved.

Before you modify an existing Virtual Chassis that uses a member MAC address as its device ID, we recommend converting it to use a virtual device ID. This makes managing the Virtual Chassis easier. For more information, refer to "Convert a Virtual Chassis to Use a Virtual Device ID" on page 226.

Convert a Virtual Chassis to Use a Virtual Device ID

When a Virtual Chassis device is represented in Mist by the MAC address of one of its member switches, managing it can become challenging. Especially, replacing or removing a member switch may cause inconsistencies in how the Virtual Chassis is represented, potentially disrupting connectivity.

Therefore, we recommend converting any existing Virtual Chassis device that uses the member 0 MAC address as its device ID to use a virtual device ID instead. Moving to a virtual device ID provides a consistent and centralized way to represent and manage a Virtual Chassis as a single logical entity, making future operations cleaner and more reliable.

A virtual device ID starts with the value 0200. A device ID that starts with any other value is assumed to be based on a member MAC address.

You can convert a Virtual Chassis using the Mist portal or via API.

To convert a virtual chassis using the Mist portal:

1. Click **Switches** to go to the list of switches.

2. From the list of switches, identify the Virtual Chassis device which you want to convert.

3. Click the Virtual Chassis device to open it.

   The Virtual Chassis details page (switch details page) appears.

4. Click the **Modify Virtual Chassis** button on the upper right of the page.

   For a Virtual Chassis device that is using a MAC address-based device ID, the Modify Virtual Chassis page displays a banner message recommending that you convert the Virtual Chassis to use a virtual device ID instead.

5. Click the **How do I convert?** button on the right of the banner message.

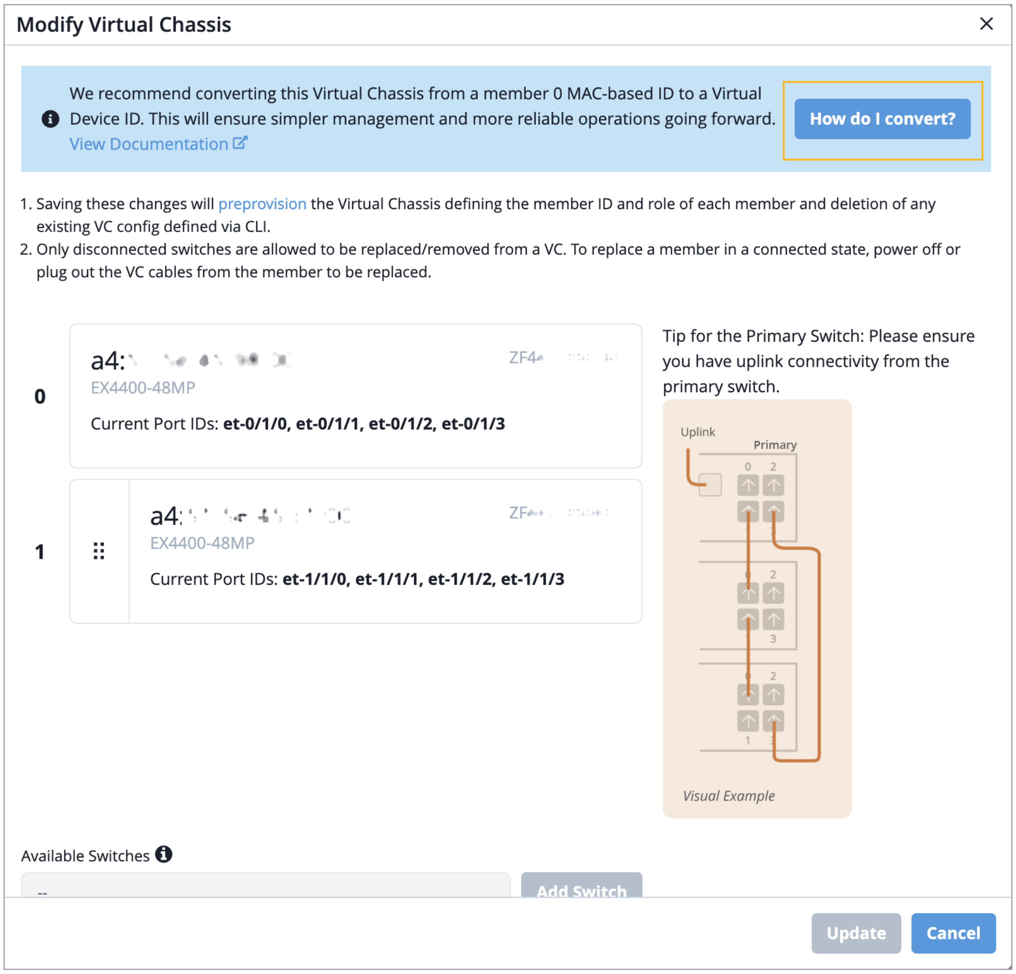

Modify Virtual Chassis

We recommend converting this Virtual Chassis from a member 0 MAC-based ID to a Virtual Device ID. This will ensure simpler management and more reliable operations going forward. View Documentation

How do I convert?

1. Saving these changes will preprovision the Virtual Chassis defining the member ID and role of each member and deletion of any existing VC config defined via CLI.
2. Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or plug out the VC cables from the member to be replaced.

0 a4:[redacted] ZF4[redacted]
EX4400-48MP
Current Port IDs: **et-0/1/0, et-0/1/1, et-0/1/2, et-0/1/3**

1 a4:[redacted] ZF4[redacted]
EX4400-48MP
Current Port IDs: **et-1/1/0, et-1/1/1, et-1/1/2, et-1/1/3**

Tip for the Primary Switch: Please ensure you have uplink connectivity from the primary switch.

Available Switches

--

Add Switch

Update

Cancel

6. Click **Convert** to initiate the conversion.

After conversion, you will be redirected to the switch list page. Following conversion, the existing Virtual Chassis device (represented by a member MAC) is disconnected from the Mist cloud, and a new device along with a virtual device ID is created and displayed. A virtual device ID typically starts with '0200'.

To convert a virtual chassis via API:

1. Locate the site where the Virtual Chassis is deployed and identify the site ID. To do that, use the steps below:

   a. Navigate to **Organization** > **Site Configuration**.

   b. Select the site to open it. You can find the site ID on the **Information** tile on this page.

2. Identify the current device ID of the Virtual Chassis that is represented using a member MAC address. To do that, use the steps below:

   a. Click **Switches** to navigate to the switches page and locate the Virtual Chassis that needs to be converted.

b. Click the Virtual Chassis device to open the Virtual Chassis (switch) details page. Look for the device ID in the URL. A MAC address-based device ID typically starts with a value other than '0200'.

3. Perform the conversion by issuing a POST request to the API endpoint below using your Site ID and Device ID.

   Endpoint:

   ```
   https://api.<cloud_env>.mist.com/api/v1/sites/<site_id>/devices/<device_id>/vc/
   convert_to_virtualmac
   ```

   Example:

   ```
   POST https://api.mist.com/api/v1/sites/978c48e6-6ef6-11e6-8bbf-02e208b2d34f/devices/
   00000000-0000-0000-1000-a4e11a000000/vc/convert_to_virtualmac
   ```

Following conversion, the existing Virtual Chassis device (represented by a member MAC) is disconnected from the Mist cloud, and a new device along with a virtual device ID is created and displayed. A virtual device ID typically starts with '0200'.

NOTE:

- Converting a Virtual Chassis device to use a virtual device ID will permanently erase all the events and stats previously stored for this device in the Mist cloud.
- The newly created Virtual Chassis may initially show as disconnected on the switch list page. However, within a few minutes, it will automatically reconnect and appear on the switch list as connected.
- The conversion does not impact the data plane—switching functionality continues uninterrupted.

Replace a Virtual Chassis Member

NOTE: Instructions in this topic apply to any Virtual Chassis device that uses a virtual device ID, represented by a device ID starting with '0200', as shown in the image below. If the Virtual Chassis uses a member MAC address as its device identifier, you must follow the instructions in "Replace a Member Whose MAC Address is Used as Virtual

Chassis Device ID" on page 232 to avoid any connectivity disruptions during the process. You can find the device ID on the switch details page (Virtual Chassis page).

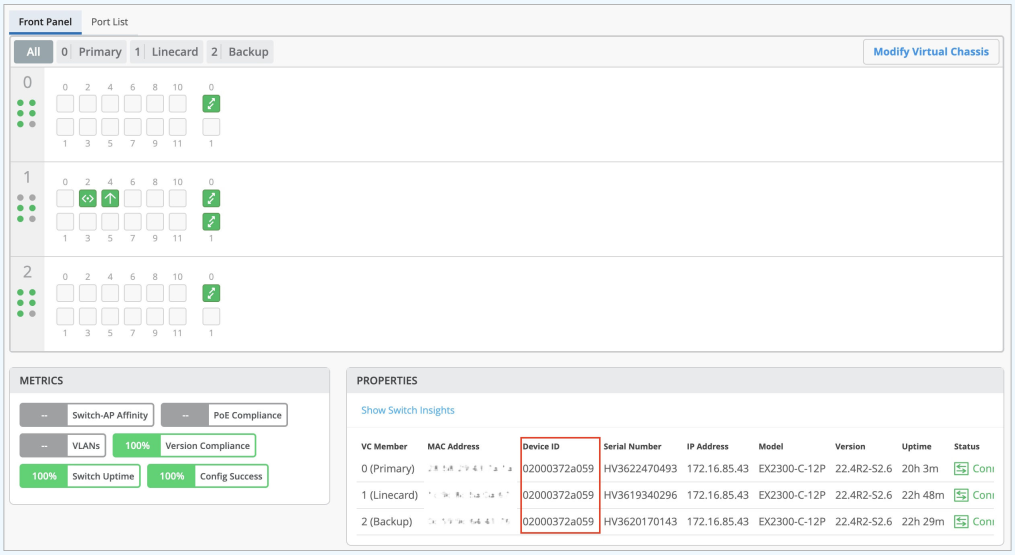

--	Switch-AP Affinity
--	PoE Compliance
--	VLANs
100%	Version Compliance
100%	Switch Uptime
100%	Config Success

VC Member	MAC Address	Device ID	Serial Number	IP Address	Model	Version	Uptime	Status
0 (Primary)	[redacted]	02000372a059	HV3622470493	172.16.85.43	EX2300-C-12P	22.4R2-S2.6	20h 3m	Con
1 (Linecard)	[redacted]	02000372a059	HV3619340296	172.16.85.43	EX2300-C-12P	22.4R2-S2.6	22h 48m	Con
2 (Backup)	[redacted]	02000372a059	HV3620170143	172.16.85.43	EX2300-C-12P	22.4R2-S2.6	22h 29m	Con

Replacing a Virtual Chassis member switch involves deleting the old member and adding a new member. Before replacing a member switch, you must ensure that:

- The new switch is of the same model family as the other members in the Virtual Chassis.
- The new switch is connected to the Virtual Chassis.
- The new switch is assigned to the same site as the other members in the Virtual Chassis.

To replace a Virtual Chassis member that has a device ID starting with '0200':

1. Onboard the replacement switch to the Mist cloud and assign it to the same site as the other members in the Virtual Chassis. To onboard the switch, use the **Claim Switch** or **Adopt Switch** option on the Inventory page (Organization > Inventory). You can also use the Mist AI mobile app to claim a switch.

   During the switch onboarding, remember to enable the configuration management for the switch by selecting the **Manage configuration with Mist** option.

   > **NOTE**: Before your replace a switch, you must ensure that the new switch is in connected state.

   For the Adopt Switch workflow, the **Manage configuration with Mist** option is available during the site assignment step. For more information on the Adopt Switch workflow, see "Onboard a Brownfield Switch" on page 18.

For the Claim Switch workflow, the **Manage configuration with Mist** option is available on the Claim Switches and Activate Subscription page. For more information about the Claim Switch workflow, see Onboard Switches to Mist Cloud.

2. If you are replacing the primary member, i.e. Routing Engine of the Virtual Chassis, perform a graceful switchover from the primary role to the backup role.

   For more information, refer to "Initiate Routing Engine Switchover " on page 242.

3. When the backup becomes the new primary, power off the original primary member (the member to be replaced). Or remove the Virtual Chassis port (VCP) cables from this member.

4. Connect the Virtual Chassis cables from the existing Virtual Chassis members to the new replacement switch.

5. On the Mist portal, navigate to the switch (Virtual Chassis) details page by clicking **Switches** > *Switch Name*.

6. Wait for the switch details page to display the member switch to be replaced as offline, as shown below:

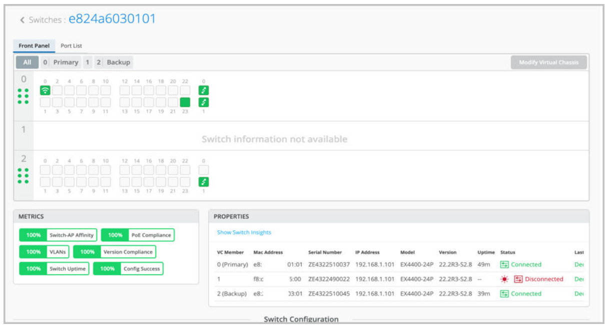

Switches : e824a6030101

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last
0 (Primary)	e8: 01:01	ZE4322510037	192.168.1.101	EX4400-24P	22.2R3-S2.8	49m	Connected	De
1	f8:c 5:00	ZE4322490022	192.168.1.101	EX4400-24P	22.2R3-S2.8	--	Disconnected	De
2 (Backup)	e8: 03:01	ZE4322510045	192.168.1.101	EX4400-24P	22.2R3-S2.8	39m	Connected	De

7. Click **Modify Virtual Chassis**.

   Because you removed the VCP connection from the member switch being replaced, the Modify Virtual Chassis window displays a broken link for this member switch along with a delete (trash) icon.

8. Delete the member to be replaced by clicking the trash icon.

9. Click **Add Switch** to add the new replacement member.

10. Renumber the new switch by dragging and dropping it into the appropriate slot.

11. Edit the MAC address of the primary or backup switch if you are replacing one of them.

12. Click **Update**.

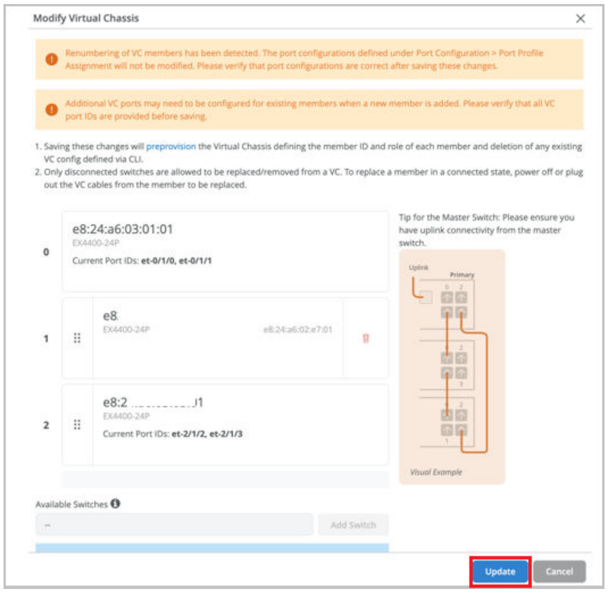

Modify Virtual Chassis

Renumbering of VC members has been detected. The port configurations defined under Port Configuration > Port Profile Assignment will not be modified. Please verify that port configurations are correct after saving these changes.

Additional VC ports may need to be configured for existing members when a new member is added. Please verify that all VC port IDs are provided before saving.

1. Saving these changes will preprovision the Virtual Chassis defining the member ID and role of each member and deletion of any existing VC config defined via CLI.
2. Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or plug out the VC cables from the member to be replaced.

0 e8:24:a6:03:01:01
EX4400-24P
Current Port IDs: **et-0/1/0, et-0/1/1**

1 e8
EX4400-24P e8:24:a6:02:e7:01

2 e8:2[redacted]1
EX4400-24P
Current Port IDs: **et-2/1/2, et-2/1/3**

Tip for the Master Switch: Please ensure you have uplink connectivity from the master switch.

Available Switches

13. Ensure that the replacement switch, which is onboarded, is powered on.

14. Wait for Virtual Chassis formation to be complete.

    The Switch Events page displays all the Virtual Chassis update events.

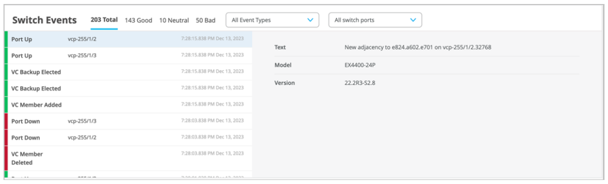

Switch Events 203 Total 143 Good 10 Neutral 50 Bad All Event Types All switch ports

Port Up	vcp-255/1/2	7:28:15.838 PM Dec 13, 2023
Port Up	vcp-255/1/3	7:28:15.838 PM Dec 13, 2023
VC Backup Elected		7:28:15.838 PM Dec 13, 2023
VC Backup Elected		7:28:15.838 PM Dec 13, 2023
VC Member Added		7:28:15.838 PM Dec 13, 2023
Port Down	vcp-255/1/3	7:28:03.838 PM Dec 13, 2023
Port Down	vcp-255/1/2	7:28:03.838 PM Dec 13, 2023
VC Member Deleted		7:28:03.838 PM Dec 13, 2023

Text	New adjacency to e824.a602.e701 on vcp-255/1/2.32768
Model	EX4400-24P
Version	22.2R3-S2.8

The switch details page displays the updated Virtual Chassis information.

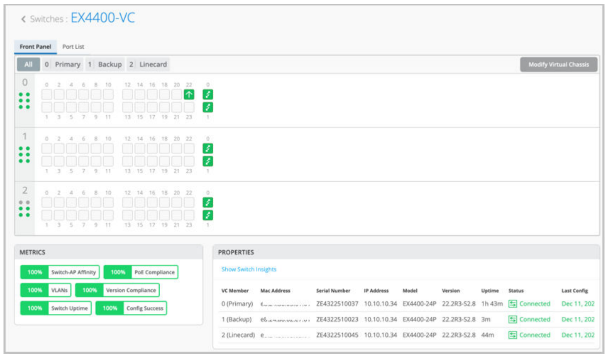

Switches : EX4400-VC

Front Panel Port List

All 0 Primary 1 Backup 2 Linecard Modify Virtual Chassis

METRICS

100% Switch-AP Affinity	100% PoE Compliance
100% VLANs	100% Version Compliance
100% Switch Uptime	100% Config Success

PROPERTIES

Show Switch Insights

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last Config
0 (Primary)	[redacted]	ZE4322510037	10.10.10.34	EX4400-24P	22.2R3-S2.8	1h 43m	Connected	Dec 11, 202
1 (Backup)	[redacted]	ZE4322510023	10.10.10.34	EX4400-24P	22.2R3-S2.8	3m	Connected	Dec 11, 202
2 (Linecard)	[redacted]	ZE4322510045	10.10.10.34	EX4400-24P	22.2R3-S2.8	44m	Connected	Dec 11, 202

Replace a Member Whose MAC Address is Used as Virtual Chassis Device ID

NOTE: Before you attempt the steps in this section, we recommend converting any existing Virtual Chassis device that uses a member MAC address as its device ID to use a virtual device ID (for more information, refer to "Convert a Virtual Chassis to Use a Virtual Device ID" on page 226). This makes managing the Virtual Chassis device easier. To replace a member in a Virtual Chassis that is converted to use a virtual device ID, you must follow the steps in "Replace a Virtual Chassis Member" on page 228.

Instructions in this topic apply to any Virtual Chassis device that uses a member MAC address as its device identifier, as shown in the image below. If the Virtual Chassis that has a device ID starting with 0200, you must follow the instructions in "Replace a Virtual Chassis Member" on page 228 to replace its members. You can find the device ID on the switch details page (Virtual Chassis page).

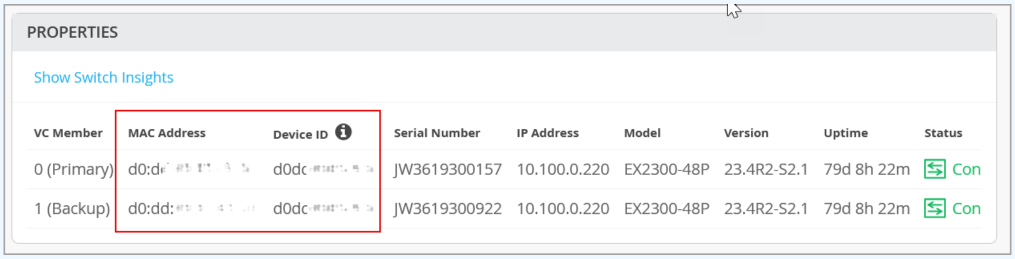

PROPERTIES

Show Switch Insights

VC Member	MAC Address	Device ID	Serial Number	IP Address	Model	Version	Uptime	Status
0 (Primary)	d0:d[redacted]	d0dc[redacted]	JW3619300157	10.100.0.220	EX2300-48P	23.4R2-S2.1	79d 8h 22m	Con
1 (Backup)	d0:dd:[redacted]	d0dc[redacted]	JW3619300922	10.100.0.220	EX2300-48P	23.4R2-S2.1	79d 8h 22m	Con

If a Virtual Chassis uses the MAC address of a member (typically the FPC0) as the device identifier, you cannot replace that member in a single operation as it is used to communicate to the Mist cloud. You need to carry out the replacement in a 2-step process that includes adding the new replacement switch and then removing the switch to be replaced. In such cases, you should carry out the member replacement operation in a maintenance window as this operation can impact the traffic to the clients connected.

NOTE: Replacing an FPC member that is used as the device identifier will freshly assign the Virtual Chassis with a new device ID that is no longer tied to any FPC member.

Before replacing a member switch, you must ensure that:

- The new switch is of the same model family as the other members in the Virtual Chassis.
- The new switch is connected to the Virtual Chassis.
- The new switch is assigned to the same site as the other members in the Virtual Chassis.

- The Virtual Chassis is pre-provisioned. For a 2-member VC, the split and merge feature is disabled by default (no-split-detection) if the Virtual Chassis is provisioned by the cloud.

To replace a member (FPC0, for example) whose MAC address is used as the Virtual Chassis device ID:

1. Remove the uplink connection (in-band or OOB) from the FPC0 member (if an uplink is present). Ensure the connectivity to the Mist cloud is maintained after the removal of the uplink. If this is the only uplink, connect it to another member that can provide the uplink connectivity.
2. If the FPC0 is a primary member, i.e. Routing Engine of the Virtual Chassis, perform a graceful switchover from the primary role to the backup role.

   For more information, refer to "Initiate Routing Engine Switchover " on page 242.
3. Power off the FPC0 member to be replaced. Or remove the VCP cable from it.
4. Claim the new member switch, assign it to the same site, and enable configuration management on the switch.

   If this member switch belongs to a model (such as the EX2300, EX4000-8P, EX4650, or QFX5120) that does not have a dedicated VC port, you should connect it to the Mist cloud.

   For information on how to claim a switch, refer to Cloud-Ready EX and QFX Switches with Mist.
5. Connect the VCP cable to the new member in the same ports.

   The new member is added to the Virtual Chassis.

   Now, the Virtual Chassis status in Junos will look like the below:

   - fpc0 - Not present
   - fpc1 - Present
   - fpc2 - Non-provisioned because the fpc0 and fpc1 were already pre-provisioned.

   > **NOTE**: The "Unprvsnd" (Unprovisioned) status will not be displayed for switch models that do not have dedicated VCPs. Mist will automatically push the necessary VCP settings to such switches after you complete the Modify Virtual Chassis steps.

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis
Virtual Chassis ID: 19e4.0553.ff90
Virtual Chassis Mode: Enabled
                                                Mstr           Mixed Route Neighbor List
```

```
Member ID  Status   Serial No    Model          prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  NotPrsnt EZ0524AX0140
1 (FPC 1)  Prsnt    FA1024AX0329 ex4100-48p     129   Master*      N  VC
-          Unprvsnd FJ1123AV0214 ex4100-f-12p
```

6. In the Mist portal, click **Switches** > *Switch Name* to go to the switch details page of the Virtual Chassis to be modified.
7. Click **Modify Virtual Chassis**.
   The Modify Virtual Chassis window appears.
8. On the Modify Virtual Chassis window, click **Add Switch** and add the replacement switch to the Virtual Chassis as a new member.
9. Click **Update**.
10. Wait for 5 to 10 minutes for the data to synchronize and appear on the Mist portal.
11. Click **Modify Virtual Chassis** again.

    Because you removed the VCP connection from the FPC0 being replaced, the Modify Virtual Chassis window displays a broken link against this member switch along with a delete (trash) icon.
12. Delete the member to be replaced by clicking the trash icon.
    The Modify Virtual Chassis window displays a message indicating that FPC0 is required.
13. Move the FPC2 member to slot 0 (the FPC0 slot) by dragging and dropping.

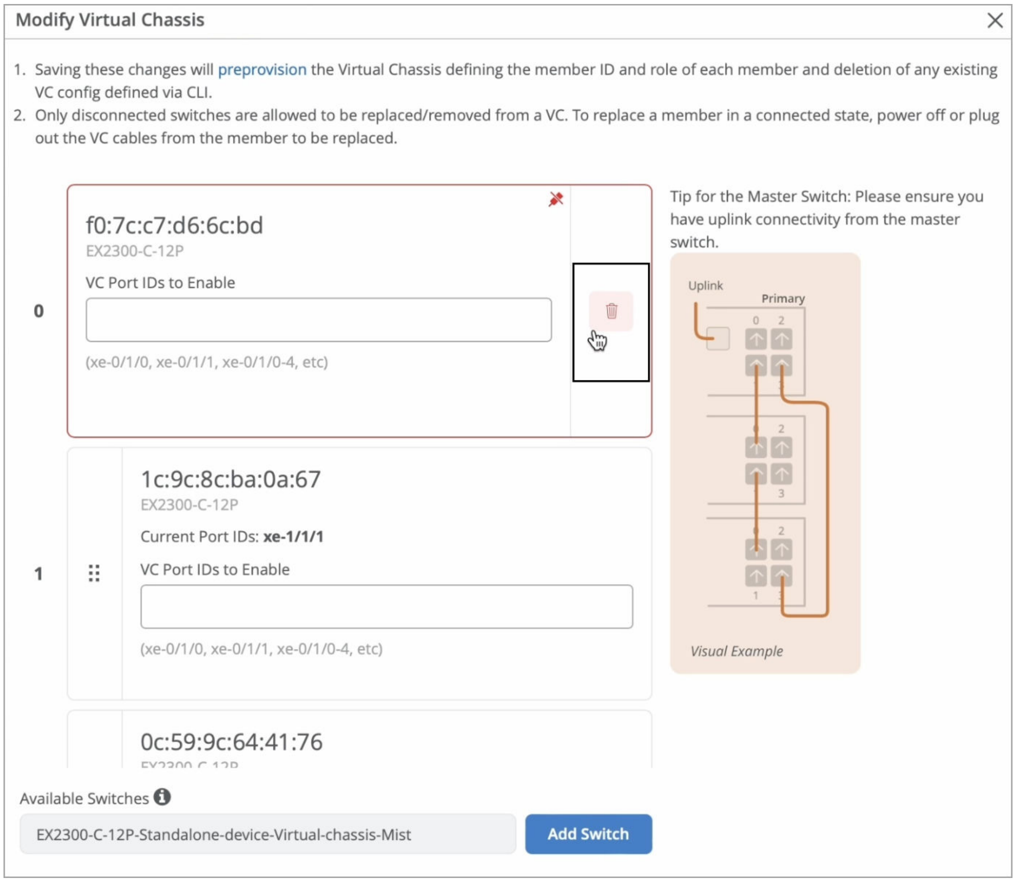

Modify Virtual Chassis ✕

1. Saving these changes will preprovision the Virtual Chassis defining the member ID and role of each member and deletion of any existing VC config defined via CLI.
2. Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or plug out the VC cables from the member to be replaced.

0

f0:7c:c7:d6:6c:bd
EX2300-C-12P
VC Port IDs to Enable

(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

1

1c:9c:8c:ba:0a:67
EX2300-C-12P
Current Port IDs: **xe-1/1/1**
VC Port IDs to Enable

(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

0c:59:9c:64:41:76
EX2300-C-12P

Tip for the Master Switch: Please ensure you have uplink connectivity from the master switch.

Available Switches

EX2300-C-12P-Standalone-device-Virtual-chassis-Mist

NOTE: Ensure that no role change is performed.

14. Click **Update**.
15. Update the **Routing Engine 2** field with the MAC address of the new switch.
16. Click **Update**.

Renumber the Virtual Chassis Members

If you prefer to see the Virtual Chassis members on the Mist portal in the same order as they are physically stacked, you need to reorder the switches (after they are powered on and connected to Virtual Chassis) using the Modify Virtual Chassis option.

You can modify the member switches' order on the Mist portal by renumbering the members. On the Modify Virtual Chassis window, accessible from the switch details page, you can move around the port panel of a switch to change the order of the member. The order is incremental. The first entry is member 0, the second is member 1, and so on. You are required to specify the FPC0.

To renumber the switches in a preprovisioned Virtual Chassis:

1. Click the **Switches** tab on the left to navigate to the Switches page.
2. Click the Virtual Chassis in which you want to renumber the members.
   The switch details page appears.
3. On the switch details page, click **Modify Virtual Chassis**.
   The Modify Virtual Chassis window appears.
4. On the Modify Virtual Chassis screen, drag and drop the port panel of a switch to different slots to change the switch number. The order is incremental. The first entry is member 0, the second is member 1, and so on. In the example below, the FPC1 has been renumbered as FPC2 and the FPC2 has been renumbered as FPC1.

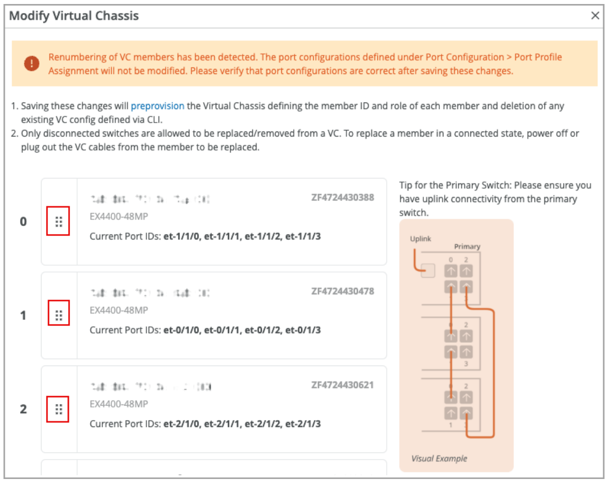

Modify Virtual Chassis

Renumbering of VC members has been detected. The port configurations defined under Port Configuration > Port Profile Assignment will not be modified. Please verify that port configurations are correct after saving these changes.

1. Saving these changes will preprovision the Virtual Chassis defining the member ID and role of each member and deletion of any existing VC config defined via CLI.
2. Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or plug out the VC cables from the member to be replaced.

0 [redacted] ZF4724430388
EX4400-48MP
Current Port IDs: **et-1/1/0, et-1/1/1, et-1/1/2, et-1/1/3**

1 [redacted] ZF4724430478
EX4400-48MP
Current Port IDs: **et-0/1/0, et-0/1/1, et-0/1/2, et-0/1/3**

2 [redacted] ZF4724430621
EX4400-48MP
Current Port IDs: **et-2/1/0, et-2/1/1, et-2/1/2, et-2/1/3**

Tip for the Primary Switch: Please ensure you have uplink connectivity from the primary switch.

NOTE:

- Within a Virtual Chassis that uses the MAC address of a member (typically FPC0) as the device identifier, you cannot renumber or move around that member unless it is disconnected.

- Renumbering the members within a Virtual Chassis does not renumber the port configurations and port profile assignment. When you renumber a VC, Mist displays the following warning (as shown in the picture above): "Renumbering of VC members has been detected. The port configurations defined under Port Configuration > Port Profile Assignment will not be modified. Please verify that port configurations are correct after saving these changes."

  So, ensure that these changes are taken care of before or after renumbering the members in the Virtual Chassis.

5. After you have made the changes, click **Update**.
   The members are renumbered.

Reassign Virtual Chassis Member Roles

A Virtual Chassis configuration in a Juniper Mist™ network has two switches in the Routing Engine role - one in the primary Routing Engine role, and the other in the backup Routing Engine role. The remaining member switches operate in the linecard role.

To change the role of Virtual Chassis members:

1. Click the **Switches** tab on the left to navigate to the Switches page.
2. Click the Virtual Chassis in which you want to change the member roles.

The switch details page appears.

3. On the switch details page, click **Modify Virtual Chassis**.

   The Modify Virtual Chassis window appears.

4. On the Modify Virtual Chassis screen, specify Routing Engine 1 or Routing Engine 2. All the other switches assume a linecard role.

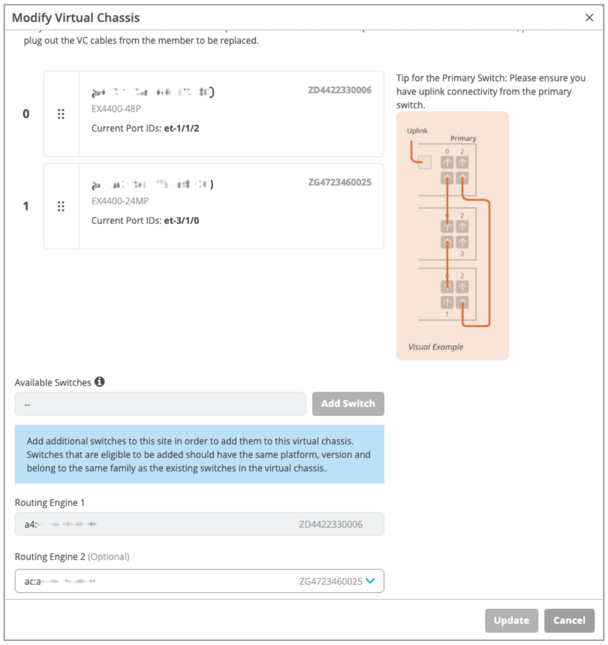

Modify Virtual Chassis

plug out the VC cables from the member to be replaced.

0 [redacted] ZD4422330006
EX4400-48P
Current Port IDs: **et-1/1/2**

1 [redacted] ZG4723460025
EX4400-24MP
Current Port IDs: **et-3/1/0**

Tip for the Primary Switch: Please ensure you have uplink connectivity from the primary switch.

Available Switches

--

Add additional switches to this site in order to add them to this virtual chassis. Switches that are eligible to be added should have the same platform, version and belong to the same family as the existing switches in the virtual chassis.

Routing Engine 1: a4:[redacted] ZD4422330006

Routing Engine 2 (Optional): ac:a[redacted] ZG4723460025

5. After you have made the changes, click **Update**.

   The member roles are changed.

You will see the updated status about the role change on the switches page on the Mist portal. The role change will take some time (approximately 15 minutes) to appear on the Mist portal. You can see a banner message at the top after every change that you make, as shown below:

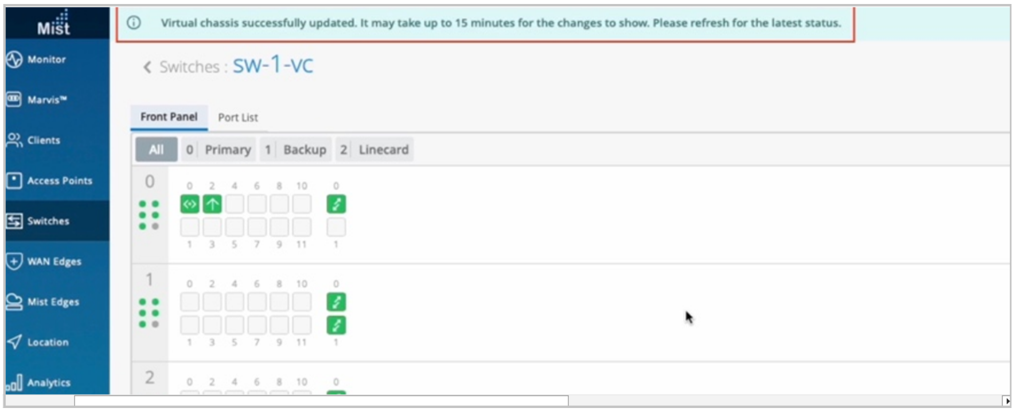

Virtual chassis successfully updated. It may take up to 15 minutes for the changes to show. Please refresh for the latest status.

Switches : SW-1-VC

Delete Virtual Chassis Members

You can delete the member switches from the Virtual Chassis, by clicking the delete (trash) icon on the Modify Virtual Chassis window. Before deleting any member switch, you must ensure that the switch to be removed is disconnected from the Virtual Chassis. If the switch is connected, power it off or remove the VCP connection from it.

To delete a member switch from Virtual Chassis:

1. If you are removing a primary member, i.e. Routing Engine of the Virtual Chassis, perform a graceful switchover from the primary role to the backup role.

   For more information, refer to "Initiate Routing Engine Switchover " on page 242.
2. Power off the member to be removed. Or remove the VCP cable from it.
3. On the Mist portal, click the **Switches** tab on the left to navigate to the Switches page.
4. Click the Virtual Chassis from which you want to delete a member switch.
   The switch details page appears.
5. On the switch details page, click **Modify Virtual Chassis**.
   The Modify Virtual Chassis window appears. Because you have removed the VCP connection from the member switch, the Modify Virtual Chassis window displays a broken link for the member switch along with a delete icon.
6. Click the delete icon.

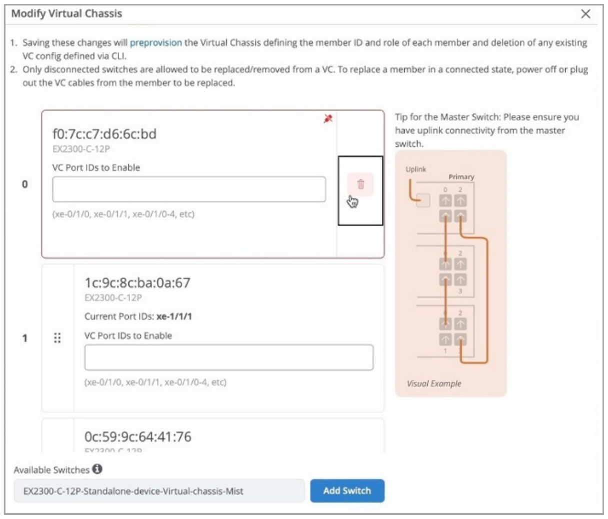

Modify Virtual Chassis

1. Saving these changes will preprovision the Virtual Chassis defining the member ID and role of each member and deletion of any existing VC config defined via CLI.
2. Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or plug out the VC cables from the member to be replaced.

0 f0:7c:c7:d6:6c:bd
EX2300-C-12P
VC Port IDs to Enable
(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

1 1c:9c:8c:ba:0a:67
EX2300-C-12P
Current Port IDs: **xe-1/1/1**
VC Port IDs to Enable
(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

0c:59:9c:64:41:76

Tip for the Master Switch: Please ensure you have uplink connectivity from the master switch.

Available Switches
EX2300-C-12P-Standalone-device-Virtual-chassis-Mist

7. Click **Update**.
   Mist removes the member switch from the Virtual Chassis.

Add a Member Switch to a Virtual Chassis

You can add one or more member switches to a Virtual Chassis from the Modify Virtual Chassis window. Before adding a new member switch to a Virtual Chassis, ensure the following:

- The new switch is of the same model family as the other members in the Virtual Chassis.
- The new switch is connected to the network (applicable to EX2300, EX4650, and QFX5120).
- The new switch is assigned to the same site as the other members in the Virtual Chassis.

> **NOTE**: Juniper Mist automatically upgrades a Virtual Chassis linecard member if it is running a Junos version different from that of the primary member. The linecard member will be upgraded to the same version as the primary member if the following conditions are met:
>
> - The switch must form a Virtual Chassis with three or more members—that is, a primary, a backup, and a linecard member.
> - The Junos version on the linecard member is different from that on the primary member.
> - The linecard member must be in Inactive state.
>
>   Note that a linecard member will be upgraded only if it is inactive and running a clearly different Junos version. Minor differences, such as different spin numbers, will not trigger an upgrade.
> - Only the Junos versions listed on the Mist portal are available for upgrade.

To add a new member switch to the Virtual Chassis:

1. Onboard the new switch to the Mist cloud and assign it to the same site as the other members in the Virtual Chassis. To onboard the switch, use the **Claim Switch** or **Adopt Switch** option on the Inventory page (Organization > Inventory). You can also use the Mist AI mobile app to claim a switch.

   During the switch onboarding, remember to enable the configuration management for the switch by selecting the **Manage configuration with Mist** option.

   For the Adopt Switch workflow, the **Manage configuration with Mist** option is available during the site assignment step. For more information on the Adopt Switch workflow, see "Onboard a Brownfield Switch" on page 18.

   For the Claim Switch workflow, the **Manage configuration with Mist** option is available on the Claim Switches and Activate Subscription page. For more information about the Claim Switch workflow, see Onboard Switches to Mist Cloud.

2. On the Mist portal, click the **Switches** tab on the left to navigate to the Switches page.
3. Click the Virtual Chassis to which you want to add the new member switch.

   The switch details page appears.
4. On the switch details page, click **Modify Virtual Chassis**.

   The Modify Virtual Chassis window appears.
5. On the Modify Virtual Chassis window, click **Add Switch**.

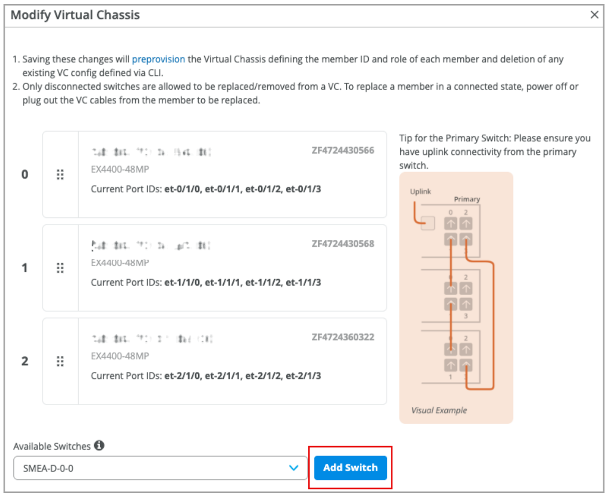

Modify Virtual Chassis

1. Saving these changes will preprovision the Virtual Chassis defining the member ID and role of each member and deletion of any existing VC config defined via CLI.
2. Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or plug out the VC cables from the member to be replaced.

0 — [redacted] ZF4724430566 EX4400-48MP Current Port IDs: **et-0/1/0, et-0/1/1, et-0/1/2, et-0/1/3**

1 — [redacted] ZF4724430568 EX4400-48MP Current Port IDs: **et-1/1/0, et-1/1/1, et-1/1/2, et-1/1/3**

2 — [redacted] ZF4724360322 EX4400-48MP Current Port IDs: **et-2/1/0, et-2/1/1, et-2/1/2, et-2/1/3**

Tip for the Primary Switch: Please ensure you have uplink connectivity from the primary switch.

6. Specify the VC port ID for the switch, if needed (the port ID configuration applies to the EX2300, EX4650, and QFX5120 switches).
7. Click **Update**.
8. Connect the VCPs as specified on the Modify Virtual Chassis window and wait for 3 to 5 minutes for virtual chassis to be updated.

   While the Virtual Chassis is forming, the switches page displays the status as 'VC Forming'.

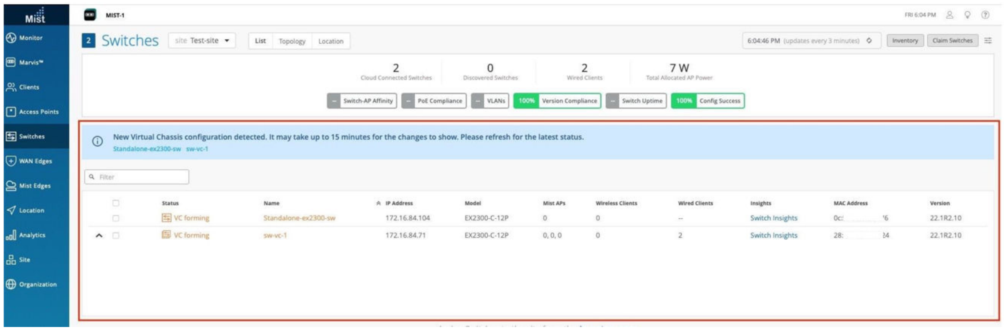

New Virtual Chassis configuration detected. It may take up to 15 minutes for the changes to show. Please refresh for the latest status.
Standalone-ex2300-sw sw-vc-1

Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	Insights	MAC Address	Version
VC forming	Standalone-ex2300-sw	172.16.84.104	EX2300-C-12P	0	0	--	Switch Insights	0c:[redacted]6	22.1R2.10
VC forming	sw-vc-1	172.16.84.71	EX2300-C-12P	0, 0, 0	0	2	Switch Insights	28:[redacted]24	22.1R2.10

After Mist updates the Virtual Chassis, the switch details page displays the front panel of all the three Virtual Chassis members.

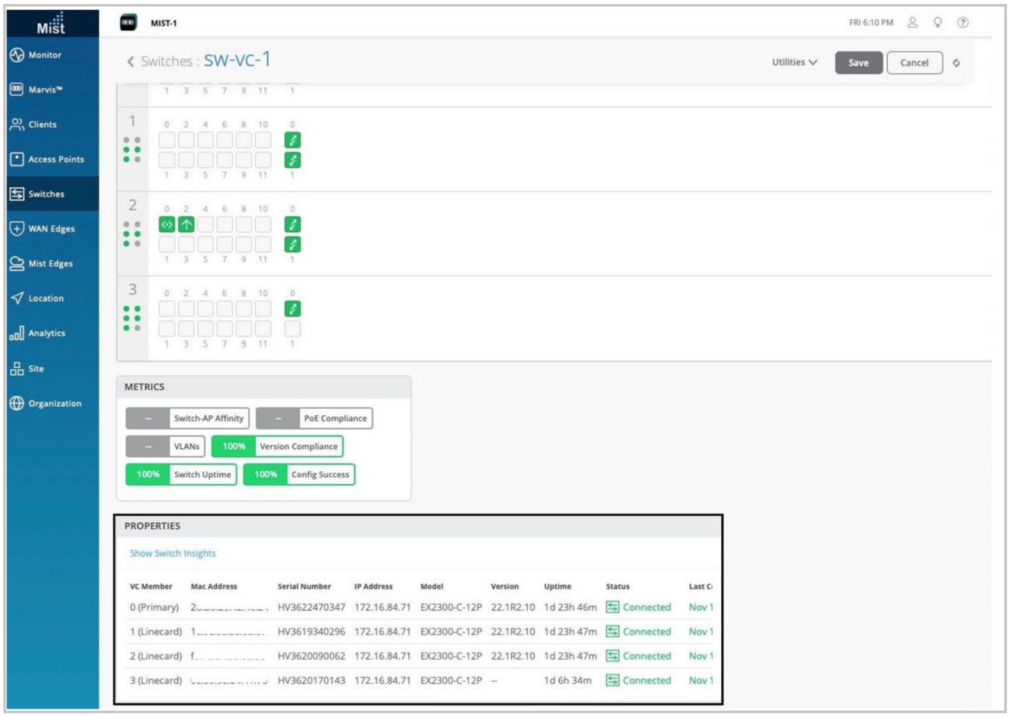

METRICS

-- Switch-AP Affinity
-- PoE Compliance
-- VLANs
100% Version Compliance
100% Switch Uptime
100% Config Success

PROPERTIES

Show Switch Insights

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last C
0 (Primary)	[redacted]	HV3622470347	172.16.84.71	EX2300-C-12P	22.1R2.10	1d 23h 46m	Connected	Nov 1
1 (Linecard)	[redacted]	HV3619340296	172.16.84.71	EX2300-C-12P	22.1R2.10	1d 23h 47m	Connected	Nov 1
2 (Linecard)	[redacted]	HV3620090062	172.16.84.71	EX2300-C-12P	22.1R2.10	1d 23h 47m	Connected	Nov 1
3 (Linecard)	[redacted]	HV3620170143	172.16.84.71	EX2300-C-12P	--	1d 6h 34m	Connected	Nov 1

Preprovision a Virtual Chassis

Before modifying any Virtual Chassis, we recommend that you ensure it is preprovisioned.

The preprovisioned configuration specifies the chassis serial number, member ID, and role for the member switches in a Virtual Chassis. When a new member router joins the Virtual Chassis, Junos compares its serial number against the values specified in the preprovisioned configuration. Preprovisioning prevents any accidental role assignments, or the accidental addition of a new member to the Virtual Chassis. Each role, member ID, addition or removal of members, is under the control of the configuration.

To preprovision a Virtual Chassis:

1. Navigate to the **Switches** page (switch list) and review the preprovisioning status in the **Preprovisioned VC** column. The Virtual Chassis devices that are not preprovisioned are highlighted with an 'x' mark in red.
2. Click to open the Virtual Chassis that has not been preprovisioned.
   The Virtual Chassis (switch) details page appears with a warning message indicating that the device is not yet preprovisioned.
3. Click the **Preprovision** button on the right side of the warning message to go to the Modify Virtual Chassis window.
4. On the Modify Virtual Chassis window, click **Preprovision Virtual Chassis**.

This action pushes the preprovisioned Virtual Chassis configuration to the device and overwrites the old autoprovision Virtual Chassis configuration pushed to the device during the ZTP process. This option assumes the current positioning of the members and preprovisions them as is.

Initiate Routing Engine Switchover

You can manually initiate Routing Engine switchover in a Virtual Chassis.

In a Virtual Chassis, one member switch is assigned the primary role and hosts the primary Routing Engine. Another member switch is assigned the backup role and hosts the backup Routing Engine. Routing Engine switchover enables the system to transition control from the primary Routing Engine to the backup Routing Engine.

To perform Routing Engine switchover:

1. Click the **Switches** tab on the left to navigate to the Switches page.
2. Click the Virtual Chassis in which you want to initiate a Routing Engine switchover.
   The switch details page appears.
3. On the switch details page, click **Modify Virtual Chassis**.
   The Modify Virtual Chassis window appears.
4. On the Modify Virtual Chassis screen, click **Switchover Routing Engine**.

   This option is available only if the following conditions are met:

   - The Virtual Chassis is preprovisioned.

   - The Virtual Chassis has a primary and a backup member.

   NOTE: Note that the switchover briefly affects the operation of Virtual Chassis.

5. Click **Continue**.
   Routing Engine switchover is initiated and you are redirected to the switch details page.

You will see the updated status about the Routing Engine switchover on the switch details page. This operation takes some time (approximately 15 minutes) to complete.

4

CHAPTER

Campus Fabric Configuration

SUMMARY

Use the information in this chapter to learn more about campus fabric, migrate to a campus fabric, and configure multihoming, core-distribution, and IP Clos.

IN THIS CHAPTER

Video Overview

Watch this video for a quick introduction to campus fabric and these three deployment options: multihoming, core-distribution, and IP Clos.

Video: Introducing Juniper Campus Fabric Workflow

What Do You Want to Do?

Table 21: Top Tasks

If you want to...	Use these resources:
Learn about the deployment options. *Which option is right for you?*	"Determine a Campus Fabric Topology" on page 245
Migrate your network to campus fabric. *Understand the strategy to migrate a traditional enterprise network-based architecture to a Juniper Campus Fabric EVPN-VXLAN architecture.*	"Migrate a Traditional Enterprise Network to a Juniper Campus Fabric" on page 252
Configure the settings for your chosen topology. *Depending on the chosen deployment option, use the appropriate link to get started configuring your campus fabric.*	• "Configure Campus Fabric EVPN Multihoming" on page 257 • "Configure Campus Fabric Core-Distribution" on page 265 • "Configure Campus Fabric IP Clos" on page 275

Determine a Campus Fabric Topology

IN THIS SECTION

Juniper Networks campus fabrics provide a single, standards-based Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) solution that you can deploy on any campus. You can deploy campus fabrics on a two-tier network with a collapsed core or a campus-wide system that involves multiple buildings with separate distribution and core layers.

NOTE: This topic lists multiple switch models that support the various campus fabric deployments. In the case of the QFX5130, though all variants support campus fabric, only the following variants are supported on Mist: QFX-5130-32CD, QFX-5130-48C, and QFX-5130-48CM.

You can build and manage a campus fabric using the Mist portal. This topic describes the following campus fabric topologies supported by Juniper Mist.

- EVPN Multihoming
- Campus Fabric Core-Distribution
- Campus Fabric IP Clos

Based on your specific requirements, you can build a campus fabric at the organization level or site level. An organization-level configuration is used when you want a single, unified fabric across multiple sites. Note that an organization-level topology only serves use cases where sites are connected via a common pair of cores. A site-level configuration is used when each site operates independently.

NOTE: The topology type EVPN Multihoming is available only for the site-specific campus fabric. You cannot build it at the organization level.

To help you determine which campus fabric to use, the following sections describe the use cases that each of the above topologies addresses:

EVPN Multihoming for Collapsed Core

The Juniper Networks campus fabrics EVPN multihoming solution supports a collapsed core architecture, which is a small to mid-size enterprise networking architecture. In a collapsed core model, you deploy up to two Ethernet switching platforms that are interconnected using technologies such as Virtual Router Redundancy Protocol (VRRP), Hot Standby Router Protocol (HSRP) and multichassis link aggregation group (MC-LAG). The endpoint devices include laptops, access points (APs), printers, and Internet of Things (IoT) devices. These endpoint devices plug in to the access layer using various Ethernet speeds, such as 100M, 1G, 2.5G, and 10G. The access layer switching platforms are multihomed to each collapsed core Ethernet switch in the core of the network.

The following image represents the traditional collapsed core deployment model:

Figure 8: Collapsed Core Topology

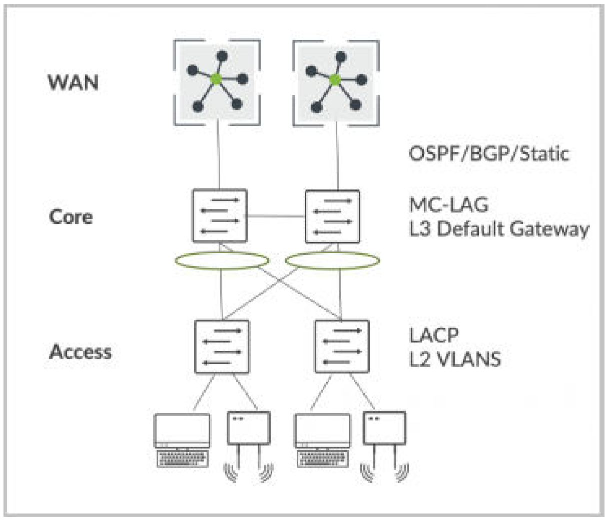

However, the traditional collapsed core deployment model presents the following challenges:

- Its proprietary MC-LAG technology requires a homogeneous vendor approach.
- It lacks horizontal scale. It supports only up to two core devices in a single topology.
- It lacks native traffic isolation capabilities in the core.
- Not all implementations support active-active load balancing to the access layer.

EVPN Multihoming addresses these challenges and provides the following advantages:

- Provides standards based EVPN-VXLAN framework.
- Supports horizontal scale up to four core devices.

- Provides traffic isolation capabilities native to EVPN-VXLAN.
- Provides native active-active load-balancing support to the access layer using Ethernet Switch Identifier-link aggregation groups (ESI-LAGs).
- Provides standard Link Aggregation Control Protocol (LACP) at the access layer.
- Mitigates the need for spanning tree protocol (STP) between the core and access layer.

Figure 9: EVPN Multihoming

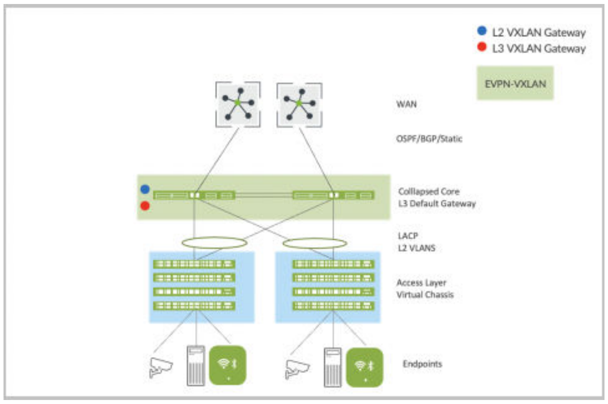

Choose EVPN Multihoming if you want to:

- Retain your investment in the access layer.
- Refresh your legacy hardware that supports collapsed core.
- Scale your deployment beyond two devices in the core.
- Leverage the existing access layer without introducing any new hardware or software models.
- Provide native active-active load-balancing support for the access layer through ESI-LAG.
- Mitigate the need for STP between the core and the access layer.
- Use the standards-based EVPN-VXLAN framework in the core.

The following Juniper platforms support EVPN Multihoming:

- Core layer devices: EX4100, EX4300-48MP, EX4400, EX4650, EX9200, QFX5120, QFX5110, QFX5700, and all QFX5130 except the QFX5130E-32CD
- Access layer devices: Third party devices using LACP, Juniper Virtual Chassis, or standalone EX switches

Campus Fabric Core-Distribution for Traditional 3-Stage Architecture

Enterprise networks that scale past the collapsed core model typically deploy a traditional three-stage architecture involving the core, distribution, and access layers. In this case, the core layer provides the Layer 2 (L2) or Layer 3 (L3) connectivity to all users, printers, APs, and so on. And the core devices interconnect with the dual WAN routers using standards-based OSPF or BGP technologies.

Figure 10: 3-Stage Core-Distribution-Access Network

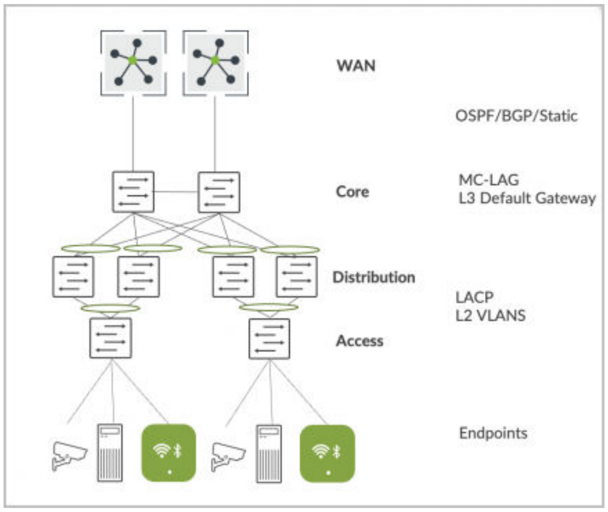

This traditional deployment model faces the following challenges:

- Its proprietary core MC-LAG technology requires a homogeneous vendor approach.
- Only up to two core devices are supported in a single topology.
- Lack of native traffic isolation capabilities anywhere in this network.
- Requires STP between the distribution and access layers and potentially between the core and distribution layers. This results in sub-optimal use of links.
- Careful planning is required if you need to move the L3 boundary between core and distribution layers.
- VLAN extensibility requires deploying VLANs across all links between access switches.

The Campus Fabric Core-Distribution architecture addresses these challenges in the physical layout of a three-stage model and provides the following advantages:

- Helps in retaining your investment in the access layer. In an enterprise network, your company makes most of the Ethernet switching hardware investment in the access layer where endpoints terminate.

The endpoint devices (including laptops, APs, printers, and IOT devices) plug in to the access layer. These devices use various Ethernet speeds, such as 100M, 1G, 2.5G, and 10G.

- Provides a standards-based EVPN-VXLAN framework.
- Supports horizontal scale at the core and distribution layers, supporting an IP Clos architecture.
- Provides traffic isolation capabilities native to EVPN-VXLAN.
- Provides native active-active load balancing to the access layer using ESI-LAG.
- Provides standard LACP at the access layer.
- Mitigates the need for STP between all layers.
- Supports the following topology subtypes:
  - Centrally routed bridging (CRB): Targets north-south traffic patterns with the L3 boundary or default gateway shared between all core devices.
  - Edge-routed bridging (ERB): Targets east-west traffic patterns and IP multicast with the L3 boundary or the default gateway shared between all distribution devices.

To know about more benefits of Campus Fabric Core-Distribution deployments, See Benefits of Campus Fabric Core-Distribution.

Figure 11: Campus Fabric Core-Distribution - CRB or ERB

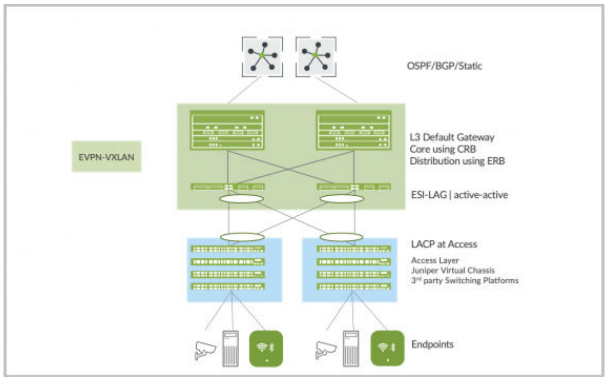

Choose Campus Fabric Core-Distribution if you want to:

- Retain your investment in the access layer while leveraging the existing LACP technology.
- Retain your investment in the core and distribution layers.
- Have an IP Clos architecture between core and distribution built on standards-based EVPN-VXLAN.

- Have active-active load-balancing at all layers, as listed below:
  - Equal-cost multipath (ECMP) between the core and distribution layers
  - ESI-LAG towards the access layer
- Mitigate the need for STP between all layers.

The following Juniper platforms support Campus Fabric Core-Distribution (CRB/ERB):

- Core layer devices: EX4650, EX9200, EX4400-48F, EX4400-24X, QFX5120, QFX5110, QFX5700, and QFX5130
- Distribution layer devices: EX4650, EX9200, EX4400-48F, EX4400-24X, QFX5120, QFX5110, QFX5700, and QFX5130
- Access layer devices: Third party devices using LACP, Juniper Virtual Chassis, or standalone EX switches

Campus Fabric IP Clos for Micro-Segmentation at Access Layer

Enterprise networks need to accommodate the growing demand for cloud-ready, scalable, and efficient networks. This demand includes a great number of IoT and mobile devices. This also creates the need for segmentation and security. IP Clos architectures help enterprises meet these challenges. An IP Clos solution provides increased scalability and segmentation using a standards-based EVPN-VXLAN architecture with Group Based Policy (GBP) capacity.

A Campus Fabric IP Clos architecture provides the following advantages:

- Micro-segmentation at the access layer using standards-based Group Based Policy
- Integration with third-party network access control (NAC) or RADIUS deployments
- Standards-based EVPN-VXLAN framework across all layers
- Flexibility in scale supporting 3-stage and 5-stage IP Clos deployments

> **NOTE**: The IP Clos architecture also supports a two-stage topology consisting of an access layer and a core layer, with the core layer acting as the services block.

- Traffic-isolation capabilities native to EVPN-VXLAN
- Native active-active load balancing within campus fabric by utilizing ECMP

- Network optimized for IP multicast
- Fast convergence between all layers, using a fine-tuned Bidirectional Forwarding Detection (BFD)
- Optional Services Block for customers who wish to deploy a lean core layer
- Mitigated need for STP between all layers

To know about more benefits of Campus Fabric IP Clos deployments, see Benefits of Campus Fabric IP Clos.

The following images represents the 3-stage and 5-stage IP Clos deployment.

Figure 12: Campus Fabric IP Clos 3 Stage

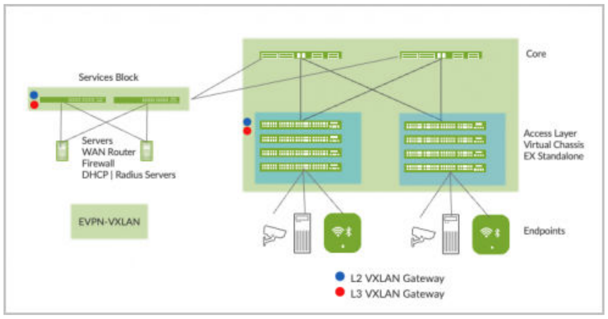

Figure 13: Campus Fabric IP Clos 5 Stage

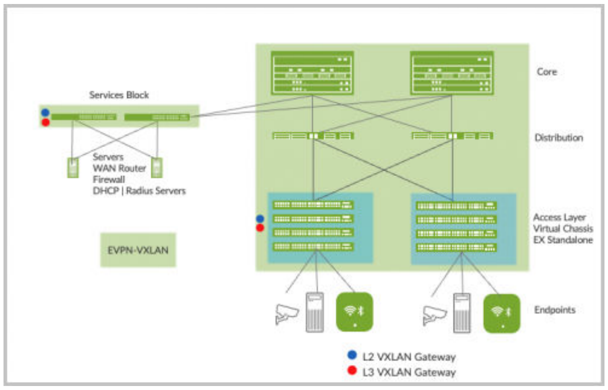

The following Juniper Network platforms support Campus Fabric IP Clos:

- Core layer devices: EX9200, EX4400-48F, EX4400-24X, EX4650, QFX5120, QFX5110, QFX5700, and QFX5130
- Distribution layer devices: EX9200, EX4400-48F, EX4400-24X, EX4650, QFX5120, QFX5110, QFX5700, and QFX5130

- Access layer devices: EX4100, EX4300-MP, and EX4400
- Services Block devices: QFX5120, EX4650, EX4400-24X, EX4400, QFX5130, QFX5700, EX9200, and QFX10k

Migrate a Traditional Enterprise Network to a Juniper Campus Fabric

SUMMARY

Use this information (including a video, diagrams, and a high-level workflow) to learn about migrating to a Campus Fabric architecture.

IN THIS SECTION

This document details a strategy that can be used to migrate a traditional enterprise network-based architecture to a Juniper Campus Fabric EVPN-VXLAN architecture.

Juniper's campus fabric leverages EVPN VXLAN as the underlying technology for small, mid, and large enterprise deployments. You can build and manage campus fabric by using Mist's Wired Assurance Cloud-ready framework. For additional information on Juniper's Campus Fabric, see Juniper Mist Wired Assurance datasheet.

Video: Three Step Campus Fabric

This migration strategy focuses on an enterprise network consisting of the traditional 3-stage architecture of access, distribution, and core. In this example, core provides layer 3 connectivity to all users, printers, access points (APs), and so on. And the core layer interconnects with dual WAN routers using standards based OSPF or BGP technologies.

Figure 14: Traditional Enterprise Network

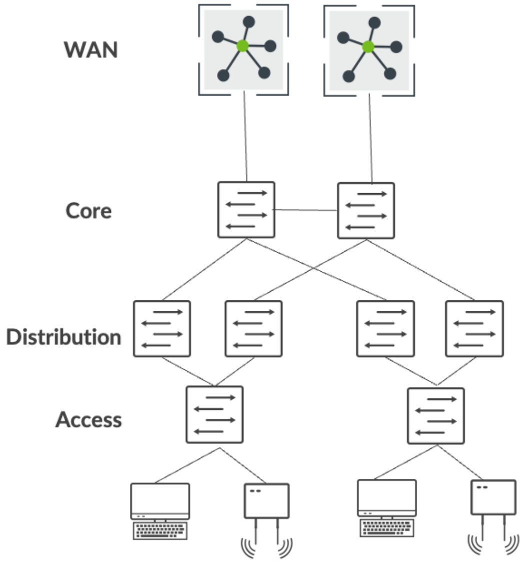

At a high-level, migration from a traditional enterprise network to a Juniper campus fabric architecture involves the following steps:

1. Build a campus fabric architecture in parallel to the existing enterprise network.

2. Interconnect the campus fabric to the existing network using a services block.

3. Migrate VLANs one by one to the campus fabric.

4. Migrate the critical infrastructure such as DHCP server and RADIUS to the services block.

5. Migrate WAN router(s) to the services block.

6. Decommission the existing enterprise network once all the connectivity is verified.

Build Campus Fabric in Parallel to the Existing Network

As the first step, build a campus fabric by using Mist's Wired Assurance framework. This step allows you to deploy an operational campus fabric in parallel to the existing network. In this example, we choose the campus fabric IP Clos architecture because the customer has a micro-segmentation strategy deployed at the access layer. The customer has chosen the following Juniper equipment to be deployed within the Campus Fabric IP Clos architecture:

- QFX5120 switches (core layer)

- QFX5120 switches (distribution layer)
- EX4100 and EX4400 switches in Virtual Chassis mode (access layer)
- QFX5120 switches (services block)

See also: "Configure Campus Fabric IP Clos" on page 275.

Figure 15: Co-existence of Campus Fabric with Enterprise Network

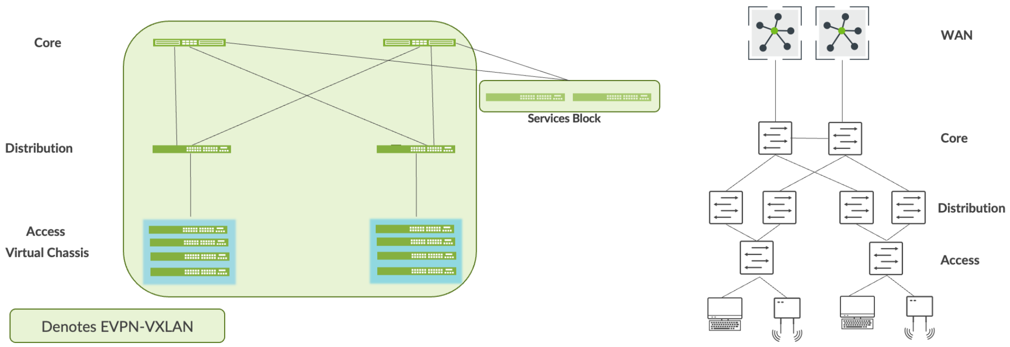

Interconnect the Campus Fabric to the Existing Network

You can use the services block to interconnect the campus fabric with the enterprise network. You can do this by using ESI-LAG technologies at layer 2, or the standard routing protocols such as BGP or OSPF if layer 3 connectivity is required. In this case, we interconnect the services block to the core enterprise using OSPF.

Figure 16: Services Block Interconnects with the Core Using OSPF

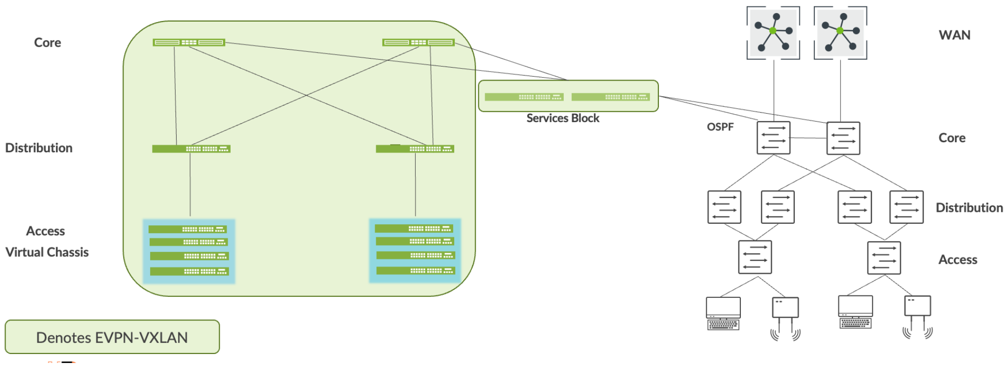

The loopback reachability between the two networks should be established through the services block. For example, the campus fabric build assigns loopback addresses to each device. By default, these addresses are part of the same subnet. OSPF should exchange these addresses with routable prefixes sent by the core layer through the services block. The end-user should verify reachability between these prefixes before moving to the next step.

Migrate VLANs to the Campus Fabric

This process requires you to remove each VLAN and associated layer 3 interface from the enterprise network. You need to migrate all devices within the VLAN to the campus fabric and then have the end-user verify full connectivity from the devices on the migrated VLAN to the applications and devices on the enterprise network. The following summarizes this step:

- Migrate VLANs to campus fabric by disabling or removing the layer 3 subnet from the current network.
- Users and devices migrate to the access layer of the campus fabric.
- Layer 3 interconnect provides reachability on a VLAN-by-VLAN basis.
- Users and devices must validate all application reachability before moving to the next VLAN.

Figure 17: All VLANs and Access Devices Migrated to the Campus Fabric

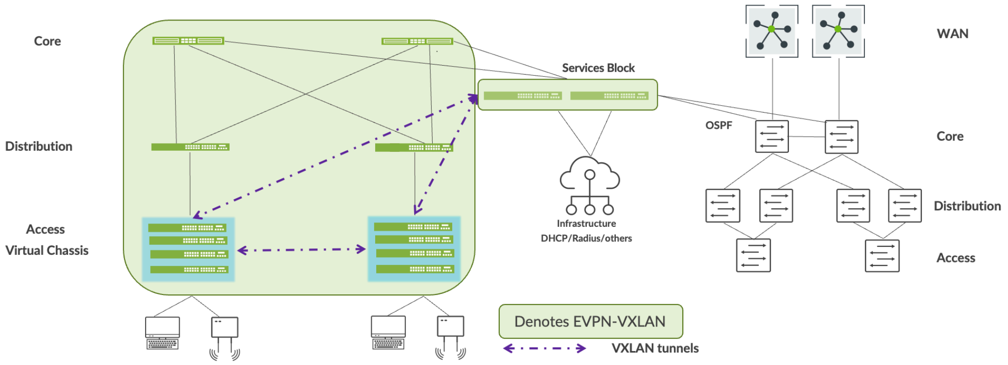

Migrate the Critical Infrastructure to the Services Block

Juniper recommends dual homing of each critical infrastructure service (such as DHCP server and RADIUS) to the services block. You can do this by using ESI-LAG technologies at layer 2, or the standard routing protocols such as BGP or OSPF if layer 3 connectivity is required. Accessibility of critical infrastructure services within the campus fabric and from the enterprise network should be verified before moving to the next step.

Figure 18: Critical Infrastructure Migration to the Services Block

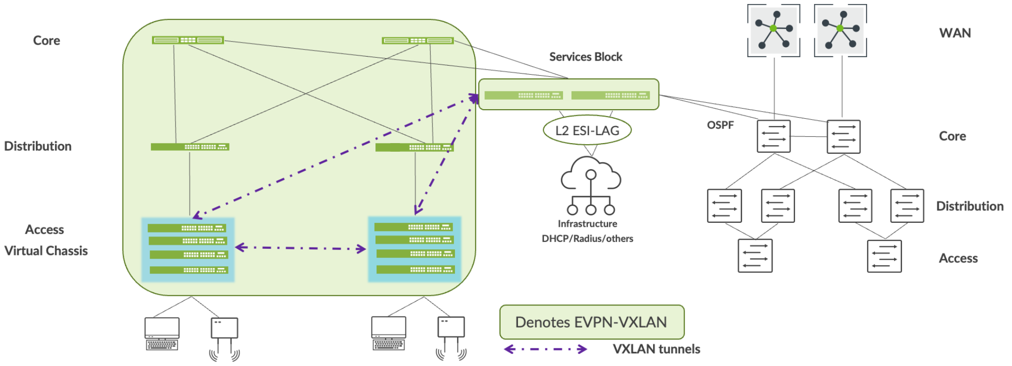

Migrate WAN Router(s) to the Services Block

Mist lets you connect WAN router(s) to the services block using BGP or OSFP. After WAN routers are connected to the services block, verify the accessibility of WAN services to and from the campus fabric before moving to the next step.

Figure 19: WAN Routers Migration to the Services Block

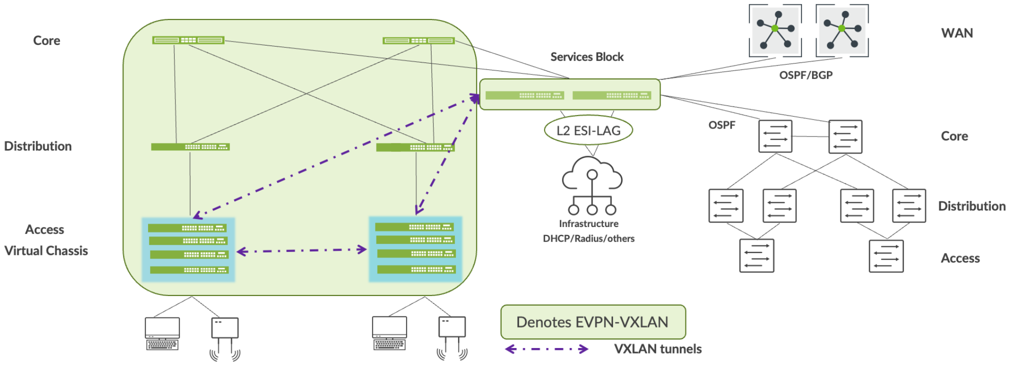

Decommission the Existing Enterprise Network

We recommend that you keep the enterprise network up and operational for at least one week after all services and applications are reachable without issue to and from the campus fabric. After that, decommission the existing enterprise network.

Configure Campus Fabric EVPN Multihoming

SUMMARY

Follow these steps to merge the core and distribution layers into a single switch by configuring EVPN multihoming.

The Juniper Networks campus fabrics EVPN multihoming solution supports a collapsed core architecture. This architecture merges the core and distribution layers into a single switch. Merging these layers into a single switch turns the traditional three-tier hierarchical network into a two-tiered network. This architecture also eliminates the need for STP across campus networks by providing multihoming capabilities from the access layer to the core layer.

NOTE:

- The topology type EVPN Multihoming is available only for the site-specific campus fabric. You cannot build it at the organization level.
- In topologies that are built in Mist cloud after the May 2025 updates, Mist automatically detects and reports any EVPN loops and duplicate MAC addresses. These issues are displayed on the switch Insights page.
  - **EVPN loop detection**—EVPN-VXLAN lightweight PE-CE loop detection helps in detecting and breaking LAN Ethernet loops on downstream leaf-to-server or access ports. This feature can detect loops caused by issues such as miswired fabric components or third-party switches incorrectly connected to the fabric.For

this feature to work, the switch must run the Junos OS version 24.4R1 or later. For more information, refer to EVPN-VXLAN Lightweight Leaf to Server Loop Detection.

- **Duplicate MAC address detection**—Identifies and mitigates issues arising from MAC address movement (MAC mobility) between different interfaces or devices in EVPN environments. While some MAC mobility is expected (for example, when a device actually moves), rapid changes might indicate issues such as network loops or misconfigurations. For more information, refer to Configuring Loop Detection for Duplicate MAC Addresses.

For a detailed configuration example, see Campus Fabric EVPN Multihoming Workflow.

Campus Fabric Configuration Best Practices

- Configure VLANs at switch template level and import them while configuring campus fabric. Template must be the single source of truth for all VLANs and port profiles unless specifically required at the switch or site level.
- At the access layer, avoid using trunk port profiles that allow all VLANs, unless explicitly required.
- Create VRF and VRF network configuration via the campus fabric, not via switch templates.
- Create port assignments per role and overwrite the configuration on individual device as needed.
- Manage the DHCP relay configuration via the campus fabric workflow, except for the service block devices.

To configure campus fabric EVPN multihoming:

1. From the left navigation menu, select **Organization** > **Campus Fabric**.
2. From the site drop-down list beside the page heading, select the site where you want to build the campus fabric.

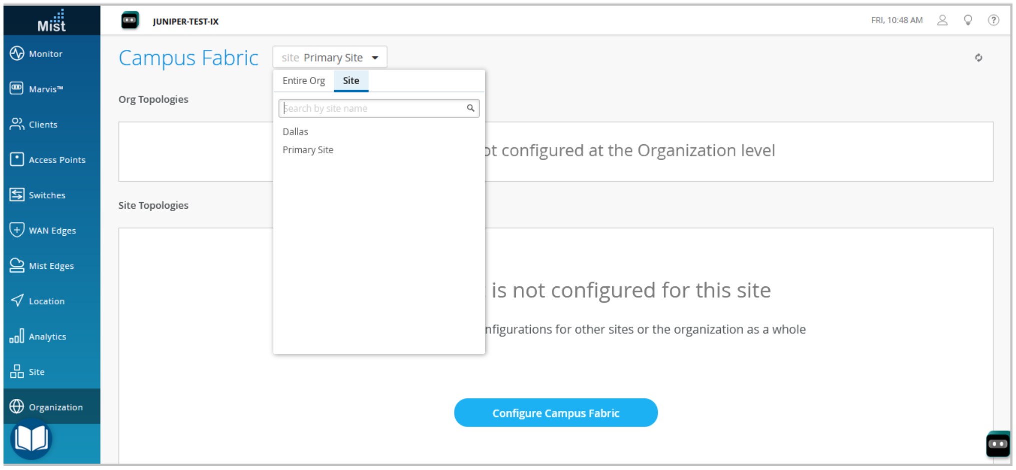

ot configured at the Organization level

is not configured for this site

nfigurations for other sites or the organization as a whole

3. Click whichever option is relevant. Click:

   - the **Configure Campus Fabric** button (displayed if the site doesn't have a campus fabric configuration associated with it).
   - the **Create Campus Fabric** button (displayed if the site already has at least one campus fabric configuration associated with it).

   The **Topology** tab is displayed.

4. Select the topology type **EVPN Multihoming**.

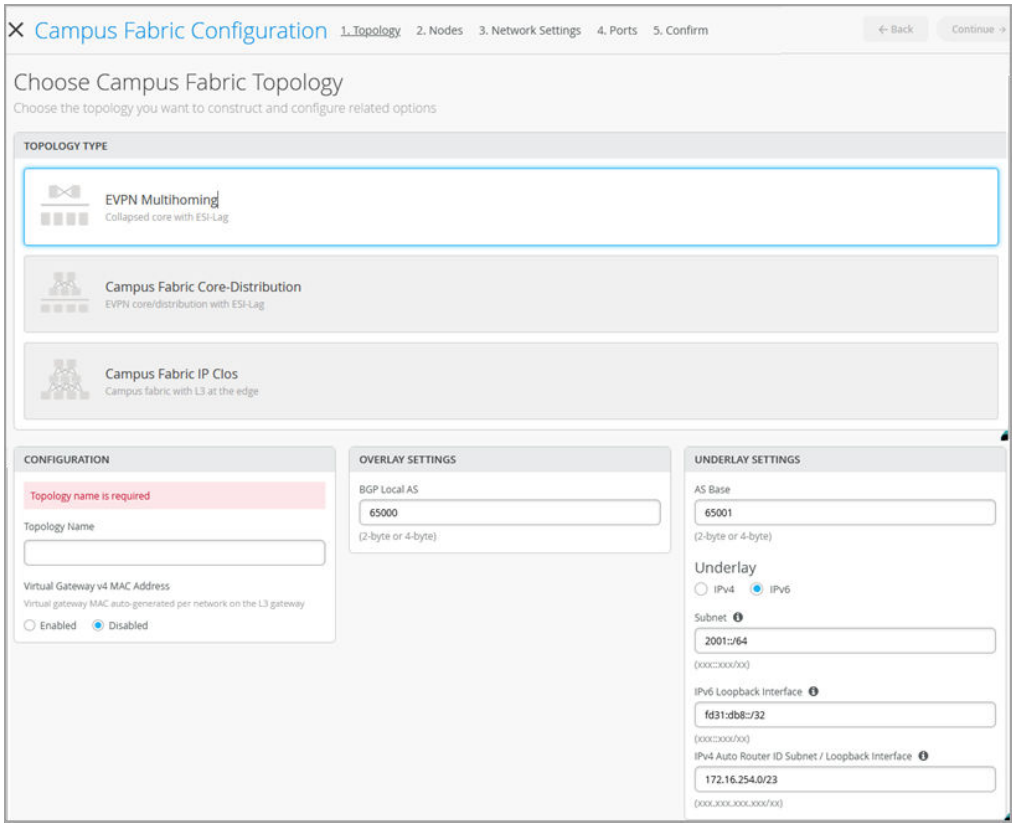

Choose Campus Fabric Topology

Choose the topology you want to construct and configure related options

TOPOLOGY TYPE

- EVPN Multihoming — Collapsed core with ESI-Lag
- Campus Fabric Core-Distribution — EVPN core/distribution with ESI-Lag
- Campus Fabric IP Clos — Campus fabric with L3 at the edge

CONFIGURATION

Topology name is required

Topology Name:

Virtual Gateway v4 MAC Address — Virtual gateway MAC auto-generated per network on the L3 gateway

○ Enabled ● Disabled

OVERLAY SETTINGS

BGP Local AS: 65000 (2-byte or 4-byte)

UNDERLAY SETTINGS

AS Base: 65001 (2-byte or 4-byte)

Underlay: ○ IPv4 ● IPv6

Subnet: 2001::/64 (xxx::xxx/xx)

IPv6 Loopback Interface: fd31:db8::/32 (xxx::xxx/xx)

IPv4 Auto Router ID Subnet / Loopback Interface: 172.16.254.0/23 (xxx.xxx.xxx.xxx/xx)

5. Configure the remaining settings on the **Topology** tab, as described below:

   > **NOTE**: We recommend that you use the default settings on this screen unless they conflict with any networks attached to the campus fabric. The point-to-point links between each layer utilize /31 addressing to conserve addresses.

   a. In the **CONFIGURATION** section, configure the following:

      - **Topology Name**—Enter a name for the topology.

- **Virtual Gateway v4 MAC Address**—Enable this option to support Guest Portal redirects on the WLAN. , Mist will provide a unique MAC address to each Layer 3 (L3) virtual gateway (per network). Disabled by default.

b. **(If you choose not to use the default settings)** In the **OVERLAY SETTINGS** section, enter the following:

- **BGP Local AS**—Represents the starting point of private BGP AS numbers that Mist allocates to each device automatically. You can use any private BGP AS number range that suits your deployment. Mist provisions the routing policy so that only the loopback IP addresses are exchanged in the underlay of the fabric.

c. **(If you choose not to use the default settings)** In the **UNDERLAY SETTINGS** section, configure the following:

- **AS Base**—The AS base number. The default is 65001.

- **Underlay**—Select an internet protocol version for the underlay. Options are IPv4 and IPv6.

- **Subnet**—The range of IP addresses that Mist uses for point-to-point links between devices. You can use a range that suits your deployment. Mist breaks this subnet into /31 subnet addressing per link. You can modify this number to suit the specific deployment scale. For example, a /24 network would provide up to 128 point-to-point /31 subnets.

- **IPv6 Loopback Interface**—Specify an IPv6 loopback interface subnet, which is used to autoconfigure IPv6 loopback interface on each device in the fabric.

- **Auto Router ID Subnet/Loopback Interface**—Mist uses this subnet to automatically assign a router ID to each device in the fabric (including access devices irrespective of whether they are configured with EVPN or not). Router IDs are loopback interfaces (lo0.0) used for overlay peering between devices. For new topologies, this field auto-populates a default subnet value (172.16.254.0/23), which you can modify. When you edit an existing topology, this field doesn’t populate any default value. The router ID is used as an identifier when deploying routing protocols such as BGP. After you add the switch to the collapsed core layer, click the switch icon to see the associated router ID.

  You can overwrite the automatically assigned router ID by manually configuring a loopback interface in the Router ID field on the Routing tile on the switch configuration page (**Switches** > ***Switch Name***). However, if you modify the campus fabric configuration afterwards, Mist performs the automatic assignment of the router ID again, replacing the manually configured loopback interface.

6. Click **Continue** to go to the **Nodes** tab, where you can select devices that form part of the campus fabric deployment.

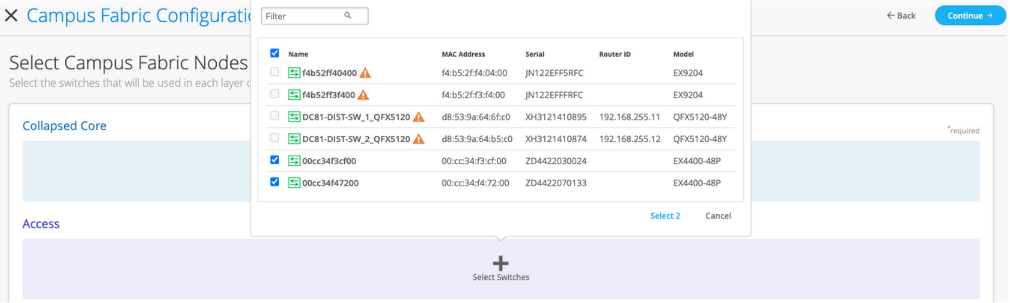

Campus Fabric Configurati

Select Campus Fabric Nodes

Select the switches that will be used in each layer c

Collapsed Core

*required

Access

Select Switches

Filter

☑	Name	MAC Address	Serial	Router ID	Model
□	f4b52ff40400 ⚠	f4:b5:2f:f4:04:00	JN122EFF5RFC		EX9204
□	f4b52ff3f400 ⚠	f4:b5:2f:f3:f4:00	JN122EFFFRFC		EX9204
□	DC81-DIST-SW_1_QFX5120 ⚠	d8:53:9a:64:6f:c0	XH3121410895	192.168.255.11	QFX5120-48Y
□	DC81-DIST-SW_2_QFX5120 ⚠	d8:53:9a:64:b5:c0	XH3121410874	192.168.255.12	QFX5120-48Y
☑	00cc34f3cf00	00:cc:34:f3:cf:00	ZD4422030024		EX4400-48P
☑	00cc34f47200	00:cc:34:f4:72:00	ZD4422070133		EX4400-48P

7. Add switches to the collapsed core layer and access layer.

   We recommend that you validate the presence of each device in the switch inventory before creating the campus fabric.

   To add switches:

   a. Click **Select Switches**.

   b. Choose the switches that you want to add to the campus fabric.

   c. Click **Select**.

8. After selecting the switches, click **Continue** to go to the **Network Settings** tab, where you can configure the networks.

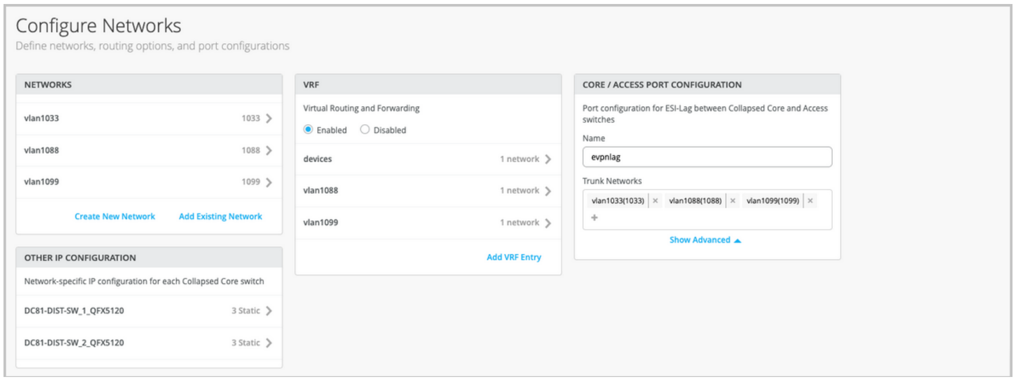

Configure Networks

Define networks, routing options, and port configurations

NETWORKS

vlan1033	1033
vlan1088	1088
vlan1099	1099

OTHER IP CONFIGURATION

Network-specific IP configuration for each Collapsed Core switch

DC81-DIST-SW_1_QFX5120	3 Static
DC81-DIST-SW_2_QFX5120	3 Static

VRF

Virtual Routing and Forwarding

● Enabled ○ Disabled

devices	1 network
vlan1088	1 network
vlan1099	1 network

CORE / ACCESS PORT CONFIGURATION

Port configuration for ESI-Lag between Collapsed Core and Access switches

Name: evpnlag

Trunk Networks: vlan1033(1033) × vlan1088(1088) × vlan1099(1099) ×

9. Configure the network settings, as described below:

   a. On the **NETWORKS** tile, add networks or VLANs to the configuration. You can either create a new network or import the network from the switch template defined on the **Organization > Switch** templates page.

      To add a new VLAN, click **Create New Network** and configure the VLANs. The settings include a name, VLAN ID, and a subnet. You can specify IPv4 or IPv6 addresses for the subnet.

      To import VLANs from the template:

      i. Click **Add Existing Network**.

      ii. Select a switch template from the **Template** drop-down list to view the VLANs available in that template.

iii. Select the required VLAN from the displayed list, and click the ✓ mark.

VLANs are mapped to Virtual Network Identifiers (VNIs). You can optionally map VLANs to virtual routing and forwarding (VRF) instances to logically separate the traffic.

b. Review the settings in the **OTHER IP CONFIGURATION** tile. This section populates the settings automatically after you specify the networks in the NETWORKS section.

c. Optionally, configure VRF instances. We recommend using VRFs when applying Type 5 (IP prefix) policies in a segment-based campus fabric architecture. By default, Mist places all VLANs in the default VRF. The VRF option allows you to group common VLANs in the same VRF or separate VRFs depending on traffic isolation requirements. All VLANs within each VRF have full connectivity with each other and with other external networking resources. A common use case is isolation of guest wireless traffic from most enterprise domains except Internet connectivity. By default, campus fabric provides complete isolation between VRFs, forcing inter-VRF communications to traverse a firewall. If you require inter-VRF communication, you need to include extra routes to the VRF. The extra route could be a default route that instructs the campus fabric to use an external router. It could also be a firewall for further security inspection or routing capabilities.

To create a VRF:

i. On the VRF tile, click **Add VRF Instance** and specify the settings. The settings include a name for the VRF and the networks to be associated with the VRF.

ii. To add extra routes, click the **Add Extra Routes** link on the **New VRF Instance** page, and specify the route. You can specify IPv4 or IPv6 addresses.

d. On the **CORE / ACCESS PORT CONFIGURATION** tile, complete the port configuration for ESI-LAG between the collapsed core and access switches. The settings include a name and other port configuration elements. By default, this configuration includes the networks added on the NETWORKS tile on the same page. If you want to remove or modify the settings, click **Show Advanced** and configure the settings. Use the tips on the screen to configure the port profile settings.

e. On the DHCP RELAY tile, configure the DHCP relay settings. You have the following options:

- Enabled—Configures DHCP relay on all the IRB-enabled devices in campus fabric. This option allows you to enable DHCP Relay on networks that you selected. The network will be populated inside the DHCP Relay tile as long as it is listed on the Networks tab on the same page.

- Disabled—Disable DHCP relay on the devices in campus fabric. When you select this option, the DHCP relay is disabled on all the IRB-enabled devices. You should carefully select this option as this will remove the locally defined DHCP Relay on the Switch Detail page.

- None—This option is automatically selected when the campus fabric topology has a mix of devices in terms of the DHCP relay configuration; that is, some devices have the DHCP relay enabled, some have it disabled, and some do not have it defined. This option will be visible for all Campus Fabric topologies that have DHCP Relay locally defined on individual switches.

  If you want to remove all locally defined DHCP Relay networks, select **Enabled** and then choose **Remove all existing device level DHCP Networks**. You can simplify your DHCP Relay deployment by centralizing any configuration change from the campus fabric workflow.

  If you enable DHCP relay in a campus fabric configuration, it is enabled on all the IRB-defined devices in the fabric and disabled on the rest of the devices. For example, in EVPN Multihoming topologies, DHCP relay is enabled on collapsed core devices and disabled on the rest.

10. Click **Continue** to go to the **Ports** tab, where you can configure the ports and create a connection between the core, distribution, and access layer switches.
11. Configure the switch ports in the collapsed core layer as follows:

    a. Select a switch in the Collapsed Core section to open the switch port panel.

    b. From the port panel of the switch, select a port that you want to configure.

    c. Specify a port type (for example, `ge` or `xe`).

    d. Select:

    - **Link to Collapsed Core** to connect the port to a core switch.

      In the collapsed core layer, you can connect each switch to every other switch to form a full mesh topology. A full mesh topology provides an EVPN Multihoming campus fabric with greater resiliency, which ensures continued network functionality even if one device fails.

    - **Link to Access** to connect the port to an access switch.

    e. Select the core or access switch (based on the selection in the previous step) on which the link should terminate. You need to configure all the ports that need to be part of the campus fabric.

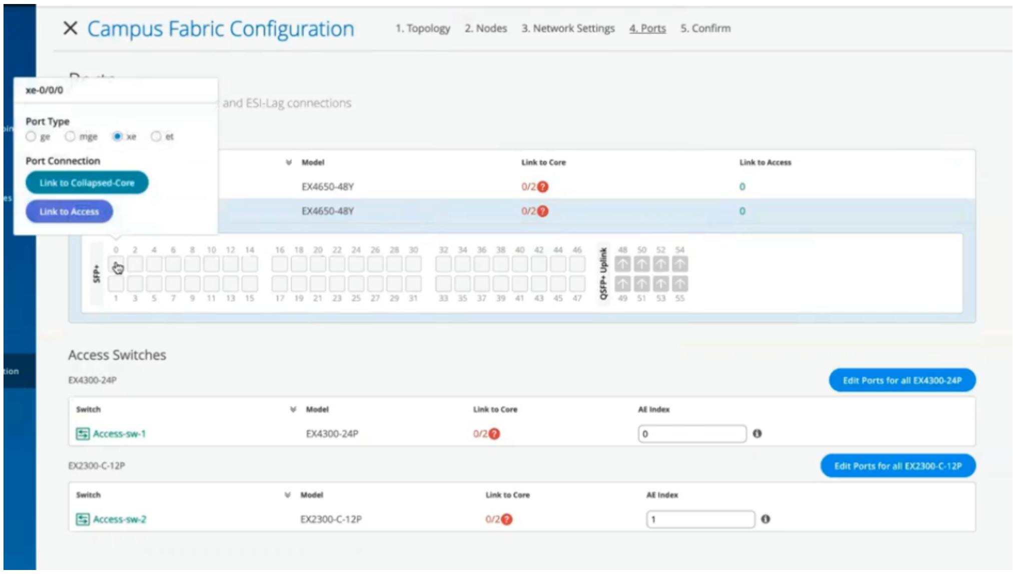

xe-0/0/0

Port Type ○ ge ○ mge ● xe ○ et

Port Connection

Link to Collapsed-Core

Link to Access

Model	Link to Core	Link to Access
EX4650-48Y	0/2	0
EX4650-48Y	0/2	0

Access Switches

EX4300-24P — Edit Ports for all EX4300-24P

Switch	Model	Link to Core	AE Index
Access-sw-1	EX4300-24P	0/2	0

EX2300-C-12P — Edit Ports for all EX2300-C-12P

Switch	Model	Link to Core	AE Index
Access-sw-2	EX2300-C-12P	0/2	1

To configure the switch ports in the access layer:

a. Select a switch in the Access section to open the switch port panel.

b. From the port panel of the switch, select a port that you want to configure.

c. Specify a port type (for example, `ge` or `xe`).

   In case the access layer uses a Virtual Chassis (VC), you can configure ports on the Primary and Backup tabs.

For the access switches, select only those interfaces that should be used to interconnect with the distribution switch. The system bundles all interfaces into a single Ethernet bundle through the AE index option. You can specify an AE index value for the access devices.

If you want to view the configuration and status information of a specific port, hover over the numbered box representing that port in the port panel UI.

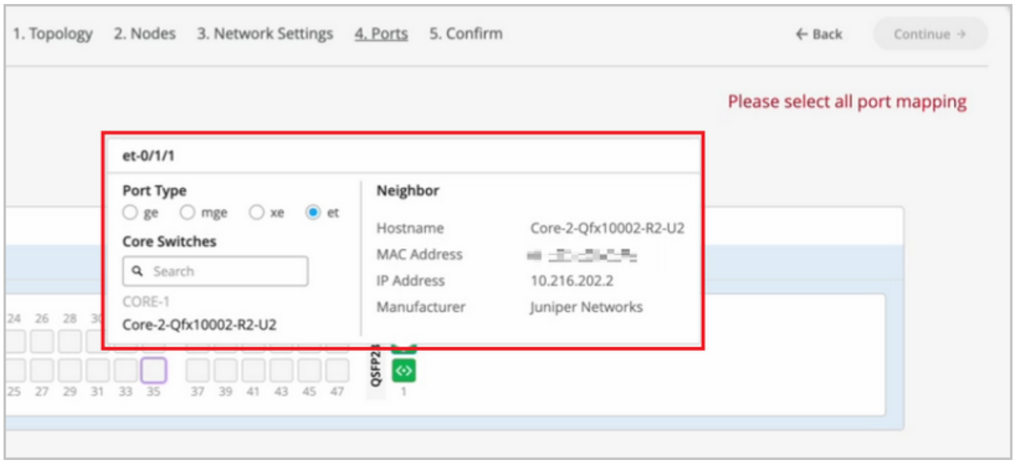

Please select all port mapping

et-0/1/1

Port Type ○ ge ○ mge ○ xe ● et

Core Switches

Search

CORE-1

Core-2-Qfx10002-R2-U2

Neighbor

Hostname	Core-2-Qfx10002-R2-U2
MAC Address	[redacted]
IP Address	10.216.202.2
Manufacturer	Juniper Networks

12. Click **Continue** to go to the **Confirmation** tab.

13. Click each switch icon to view and verify the configuration.
14. After verifying the configuration, click **Apply Changes** > **Confirm**.

    This step saves the campus fabric configuration to the Mist cloud and applies it to the switches. If the switches are offline, the configuration will be applied to them when they come online next time. A switch might take up to 10 minutes to complete the configuration.
15. Click **Close Campus Fabric Configuration**.

    After Mist builds the campus fabric, or while it is building the fabric, you can download the connection table. The connection table represents the physical layout of the campus fabric. You can use this table to validate all switch interconnects for the devices participating in the physical campus fabric build. Click **Connection Table** to download it (.csv format).

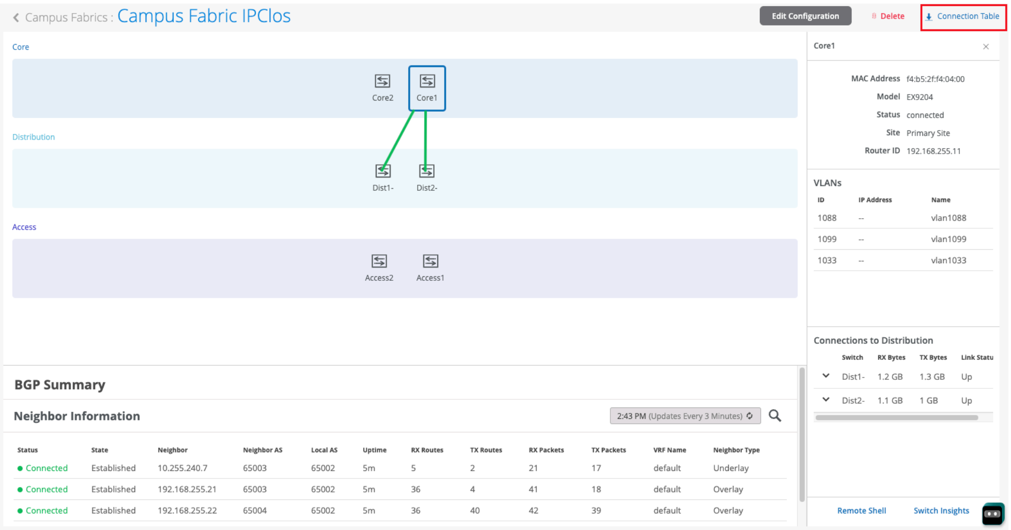

Core1

MAC Address f4:b5:2f:f4:04:00
Model EX9204
Status connected
Site Primary Site
Router ID 192.168.255.11

VLANs

ID	IP Address	Name
1088	--	vlan1088
1099	--	vlan1099
1033	--	vlan1033

Connections to Distribution

Switch	RX Bytes	TX Bytes	Link Statu
Dist1-	1.2 GB	1.3 GB	Up
Dist2-	1.1 GB	1 GB	Up

BGP Summary

Neighbor Information

2:43 PM (Updates Every 3 Minutes)

Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
● Connected	Established	10.255.240.7	65003	65002	5m	5	2	21	17	default	Underlay
● Connected	Established	192.168.255.21	65003	65002	5m	36	4	41	18	default	Overlay
● Connected	Established	192.168.255.22	65004	65002	5m	36	40	42	39	default	Overlay

16. Verify the campus fabric configuration. To verify, follow the steps listed in the **Verification** section in Campus Fabric EVPN Multihoming Workflow.

Video: Deployment of Campus Fabric EVPN Multihoming With Wired Assurance

Configure Campus Fabric Core-Distribution

SUMMARY

Follow these steps to configure a core-distribution topology for your campus fabric.

Juniper Networks campus fabrics provide a single, standards-based EVPN-VXLAN solution that you can deploy on any campus. The campus fabric core-distribution solution extends the EVPN fabric to connect VLANs across multiple buildings. This network architecture includes the core and distribution layers that integrate with the access switching layer through the standard LACP.

For more background information about campus fabric core-distribution architectures, see the following documents:

- Campus Fabric Core Distribution CRB (JVD)
- Campus Fabric Core-Distribution ERB (JVD)

NOTE:
In topologies that are built in Mist cloud after the May 2025 updates, Mist automatically detects and reports any EVPN loops and duplicate MAC addresses. These issues are displayed on the switch Insights page.

- **EVPN loop detection**—EVPN-VXLAN lightweight PE-CE loop detection helps in detecting and breaking LAN Ethernet loops on downstream leaf-to-server or access ports. This feature can detect loops caused by issues such as miswired fabric components or third-party switches incorrectly connected to the fabric.For this feature to work, the switch must run the Junos OS version 24.4R1 or later. For more information, refer to EVPN-VXLAN Lightweight Leaf to Server Loop Detection.

- **Duplicate MAC address detection**—Identifies and mitigates issues arising from MAC address movement (MAC mobility) between different interfaces or devices in EVPN environments. While some MAC mobility is expected (for example, when a device actually moves), rapid changes might indicate issues such as network loops or misconfigurations. For more information, refer to Configuring Loop Detection for Duplicate MAC Addresses.

Campus Fabric Configuration Best Practices

- Configure VLANs at switch template level and import them while configuring campus fabric. Template must be the single source of truth for all VLANs and port profiles unless specifically required at the switch or site level.

- At the access layer, avoid using trunk port profiles that allow all VLANs, unless explicitly required.

- Create VRF and VRF network configuration via the campus fabric, not via switch templates.

- Create port assignments per role and overwrite the configuration on individual device as needed.

- Manage the DHCP relay configuration via the campus fabric workflow, except for the service block devices.

To configure campus fabric core-distribution:

1. Click **Organization** > **Campus Fabric**.
2. If you want to create the campus fabric for a site, select the site from the drop-down list beside the page header. If you want to create the campus fabric for the entire organization, select **Entire Org** from the drop-down list.

   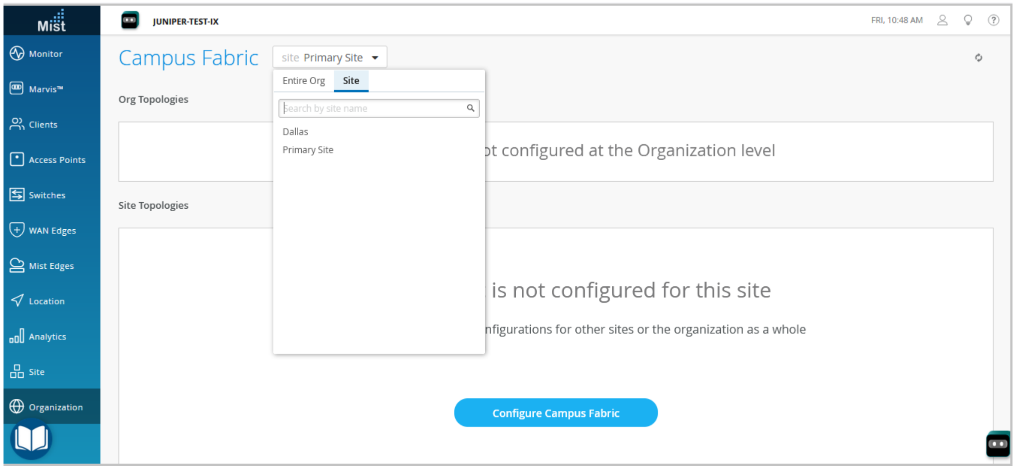

   ot configured at the Organization level

   is not configured for this site

   nfigurations for other sites or the organization as a whole

   You can use an organization-level campus fabric topology to build a campus-wide architecture with multiple buildings. Otherwise, build a site-specific campus fabric with a single set of core, distribution, and access switches.
3. Click whichever option is relevant. Click the:

   - **Configure Campus Fabric** button (displayed if the site doesn't have a campus fabric configuration associated with it).
   - **Create Campus Fabric** button (displayed if the site already has at least one campus fabric configuration associated with it).

   The **Topology** tab is displayed.
4. Select the topology type **Campus Fabric Core-Distribution**.

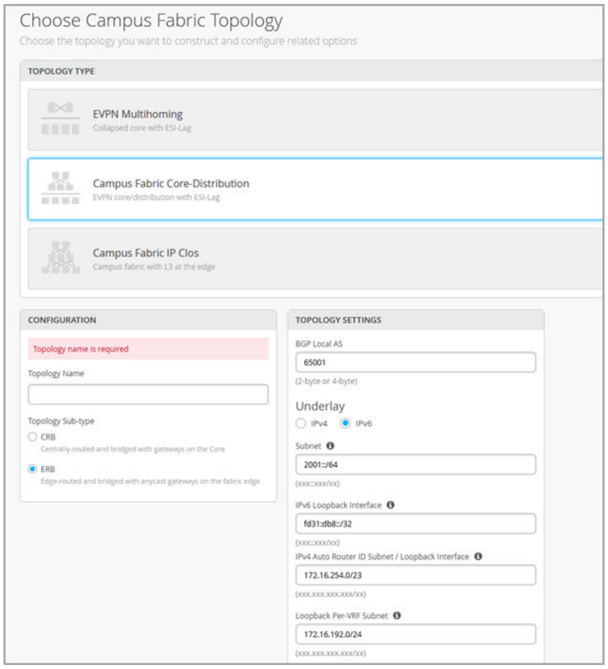

Choose Campus Fabric Topology

Choose the topology you want to construct and configure related options

TOPOLOGY TYPE

EVPN Multihoming
Collapsed core with ESI-Lag

Campus Fabric Core-Distribution
EVPN core/distribution with ESI-Lag

Campus Fabric IP Clos
Campus fabric with L3 at the edge

CONFIGURATION

Topology name is required

Topology Name

Topology Sub-type
○ CRB
Centrally-routed and bridged with gateways on the Core
● ERB
Edge-routed and bridged with anycast gateways on the fabric edge

TOPOLOGY SETTINGS

BGP Local AS
65001
(2-byte or 4-byte)

Underlay
○ IPv4 ● IPv6

Subnet
2001::/64
(xxx::xxx/xx)

IPv6 Loopback Interface
fd31:db8::/32
(xxx::xxx/xx)

IPv4 Auto Router ID Subnet / Loopback Interface
172.16.254.0/23
(xxx.xxx.xxx.xxx/xx)

Loopback Per-VRF Subnet
172.16.192.0/24
(xxx.xxx.xxx.xxx/xx)

5. Configure the topology name and other settings on the **Topology** tab, as described below:

> **NOTE**: We recommend that you use the default settings on this screen unless they conflict with any networks attached to the campus fabric. The point-to-point links between each layer utilize /31 addressing to conserve addresses.

a. In the **CONFIGURATION** section, enter the following:

- **Topology Name**—Enter a name for the topology.
- **Topology Sub-type**—Choose one of the following options:
  - **CRB**—In this model, the Layer 3 (L3) VXLAN gateway function is configured only on the core devices. This is accomplished by defining integrated routing and bridging (IRB) interfaces on the core devices to provide L3 routing services. This option uses virtual gateway addressing for all devices participating in the L3 subnet. Enabling this option configures core switches with a shared IP address for each L3 subnet. This address is shared between both the core switches and is used as the default gateway address for all devices within the VLAN. In addition, Mist assigns each core device with a unique IP address.
    - **Virtual Gateway v4 MAC Address**— Enable this option to support Guest Portal redirects on the WLAN. Mist will provide a unique MAC address to each L3 IRB interface (per network). This option is available only if you have selected CRB.
  - **ERB**—In this model, the L2 and L3 VXLAN gateway functions are configured on the distribution devices. In this case the IRB interfaces are defined on the distribution devices

to provide L3 routing services. This option uses anycast addressing for all devices participating in the L3 subnet. In this case, the distribution switches are configured with the same IP address for each L3 subnet.

b. **(If you choose not to use the default settings)** In the **TOPOLOGY SETTINGS** section, enter the following:

- **BGP Local AS**—Represents the starting point of private BGP AS numbers that Mist allocates to each device automatically. You can use any private BGP AS number range that suits your deployment. Mist provisions the routing policy so that only the loopback IP addresses are exchanged in the underlay of the fabric.

- **Underlay**—Select an internet protocol version for the underlay. Options are IPv4 and IPv6. Only ERB topologies support IPv6. You get the option to select IPv6 only if you selected ERB as Topology Sub-type.

- **Subnet**— The range of IP addresses that Mist uses for point-to-point links between devices. You can use a range that suits your deployment. Mist breaks this subnet into /31 subnet addressing per link. You can modify this number to suit the specific deployment scale. For example, a /24 network would provide up to 128 point-to-point /31 subnets.

- **IPv6 Loopback Interface**—Specify an IPv6 loopback interface subnet, which is used to autoconfigure IPv6 loopback interface on each device in the fabric.

- **IPv4 Auto Router ID Subnet / Loopback Interface**—Mist uses this subnet to automatically assign a router ID to each device in the fabric (including access devices irrespective of whether they are configured with EVPN or not). Router IDs are loopback interfaces (lo0.0) used for overlay peering between devices. For new topologies, this field auto-populates a default subnet value (172.16.254.0/23), which you can modify. When you edit an existing topology, this field doesn't populate any default value. The router ID is used as an identifier when deploying routing protocols such as BGP.

  You can overwrite the automatically assigned router ID by manually configuring a loopback interface in the Router ID field on the Routing tile on the switch configuration page (**Switches** > ***Switch Name***). However, if you modify the campus fabric configuration afterwards, Mist performs the automatic assignment of the router ID again, replacing the manually configured loopback interface.

- **Loopback per-VRF subnet**—Mist uses this subnet to automatically configure loopback interfaces (lo0.x) per the virtual routing and forwarding (VRF) instance that is used for services such as DHCP relay. For new topologies, this field auto-populates a default subnet value (172.16.192.0/24), which you can modify. This field supports a /19 or smaller subnet (for example, /24). When you edit an existing topology, this field doesn't populate any default value.

6. Click **Continue** to go to the **Nodes** tab, where you can select devices that form part of the campus fabric deployment.

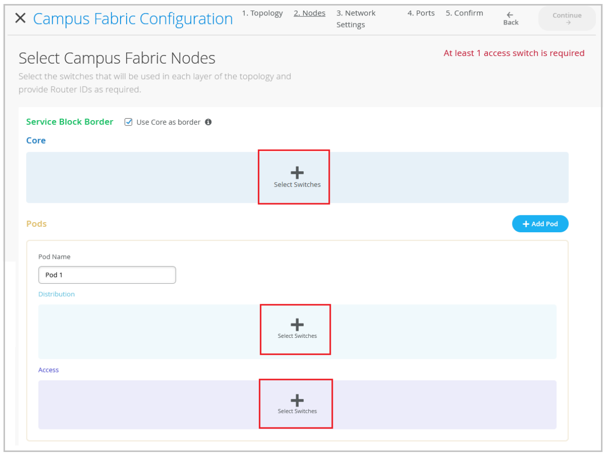

Campus Fabric Configuration

Select Campus Fabric Nodes

Select the switches that will be used in each layer of the topology and provide Router IDs as required.

At least 1 access switch is required

Service Block Border ☑ Use Core as border

Core

Pod Name: Pod 1

7. Add switches to the Core, Distribution, and Access layer sections.

   To add switches:

   a. Click **Select Switches** in the section to which you want to add switches.

   b. Choose the switches that you want to add to the campus fabric.

   c. Click **Select**.

   We recommend that you validate the presence of each device in the switch inventory before creating the campus fabric.

   By default, Mist configures the core switches to function as border nodes that run the services block functionality. In a campus fabric topology, border nodes interconnect external devices such as firewalls, routers, or critical devices. External services or devices (for example, DHCP and RADIUS servers) connect to the campus fabric through border nodes. If you want to offload this task from the core switches and use dedicated switches as border nodes, clear the **Use Core as border** checkbox on the upper left of the page. You can then add up to two switches as dedicated border nodes.

   Also, Mist provides pods for improved scalability. Your access and distribution devices are grouped into pods. A pod could represent a building. For example, you can create a pod for each of the buildings in your site and create connections between the access and the distribution devices in that pod. You do not have to connect the same set of access devices to the distribution devices across multiple buildings. You can create multiple pods by clicking **+Add Nodes**.

   You need only one connection between a pod and the core switch. You do not have to connect each distribution switch in a pod to all the core switches used. In a core-distribution topology (CRB or ERB), you need only one connection per core and distribution pair.

8. After selecting the switches, click **Continue** to go to the **Network Settings** tab, where you can configure the networks.

9. Configure the network settings, as described below:

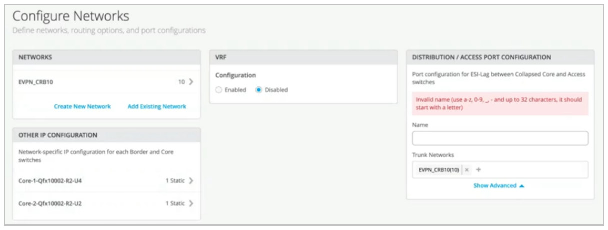

Configure Networks

Define networks, routing options, and port configurations

NETWORKS

EVPN_CRB10 10 >

OTHER IP CONFIGURATION

Network-specific IP configuration for each Border and Core switches

Core-1-Qfx10002-R2-U4 1 Static >

Core-2-Qfx10002-R2-U2 1 Static >

VRF

Configuration

○ Enabled ● Disabled

DISTRIBUTION / ACCESS PORT CONFIGURATION

Port configuration for ESI-Lag between Collapsed Core and Access switches

Invalid name (use a-z, 0-9, _, - and up to 32 characters, it should start with a letter)

Name

Trunk Networks

EVPN_CRB10(10) × +

a. On the NETWORKS tile, add networks or VLANs to the configuration. You can either create a new network or import the network from the switch template defined on the **Organization > Switch** templates page.

   To add a new VLAN, click **Create New Network** and configure the VLANs. The settings include a name, VLAN ID, and a subnet. You can specify IPv4 or IPv6 addresses for the subnet.

   You can optionally configure IPv4 and IPv6 Anycast Gateway addresses in addition to any IPv4 or IPv6 subnets you have configured. The Mist UI uses those gateways as the Anycast IP address assignment across all Access and Distribution switches in the Campus Fabric. configuration.

   > **NOTE**: If the Anycast Gateway fields are left empty, the Mist UI uses the existing logic, which is to use the first IP address in the subnet as the Anycast address.

   To import VLANs from the template:

   i. Click **Add Existing Network**.

   ii. Select a switch template from the **Template** drop-down list to view the VLANs available in that template.

   iii. Select the required VLAN from the displayed list, and click the ✓ mark.

   VLANs are mapped to Virtual Network Identifiers (VNIs). You can optionally map the VLANs to VRF instances to logically separate the traffic.

b. Review the settings on the OTHER IP CONFIGURATION tile, which populates the information automatically after you specify the networks in the NETWORKS section.

   Mist provides automatic IP addressing of IRBs for each of the VLANs. Then, the port profile associates the VLAN with the specified ports.

c. Optionally, configure VRF instances. We recommend using VRFs when applying Type 5 (IP prefix) policies in a segment-based campus fabric architecture. By default, Mist places all VLANs in the default VRF. The VRF option allows you to group common VLANs into the same VRF or separate VRFs depending on traffic isolation requirements. All VLANs within each VRF have full connectivity with each other and with other external networking resources. A common use case is the isolation of guest wireless traffic from most enterprise domains, except Internet connectivity. By default, a campus fabric provides complete isolation between VRFs, forcing inter-VRF communications to traverse a firewall. If you require inter-VRF communication, you need to include extra routes to the VRF. The extra route could be a default route that instructs the campus fabric to use an external router. It could also be a firewall for further security inspection or routing capabilities.

To create a VRF:

i. On the VRF tile, click **Add VRF Instance** and specify the settings. The settings include a name for the VRF and the networks to be associated with the VRF.

ii. To add extra routes, click the **Add Extra Routes** link on the **New VRF Instance** page and specify the route. You can specify IPv4 or IPv6 addresses.

d. On the DISTRIBUTION / ACCESS PORT CONFIGURATION tile, complete the port configuration for ESI-LAG between the collapsed core and access switches. The settings include a name and other port configuration elements. By default, this configuration includes the networks added on the NETWORKS tile on the same page. If you want to remove or modify the settings, click **Show Advanced** and configure the settings. Use the tips on the screen to configure the port profile settings.

e. On the DHCP RELAY tile, configure the DHCP relay settings. You have the following options:

- Enabled—Configures DHCP relay on all the IRB-enabled devices in campus fabric. This option allows you to enable DHCP Relay on networks that you selected. The network will be populated inside the DHCP Relay tile as long as it is listed on the Networks tab on the same page.

- Disabled—Disable DHCP relay on the devices in campus fabric. When you select this option, the DHCP relay is disabled on all the IRB-enabled devices. You should carefully select this option as this will remove the locally defined DHCP Relay on the Switch Detail page.

- None—This option is automatically selected when the campus fabric topology has a mix of devices in terms of the DHCP relay configuration; that is, some devices have the DHCP relay enabled, some have it disabled, and some do not have it defined. This option will be visible for all Campus Fabric topologies that have DHCP Relay locally defined on individual switches.

If you want to remove all locally defined DHCP Relay networks, select **Enabled** and then choose **Remove all existing device level DHCP Networks**. You can simplify your DHCP Relay deployment by centralizing any configuration change from the campus fabric workflow.

If you enable DHCP relay in a campus fabric configuration, it is enabled on all the IRB-defined devices in the fabric and disabled on the rest of the devices. For example, in Campus Fabric Core-Distribution (CRB) topologies, DHCP relay is enabled on core devices and disabled on the rest. Similarly, in Campus Fabric Core-Distribution (ERB), DHCP is enabled on distribution devices and disabled on the rest.

10. Click **Continue** to go to the **Ports** tab, where you can configure the ports and create a connection between the core, distribution, and access layer switches.

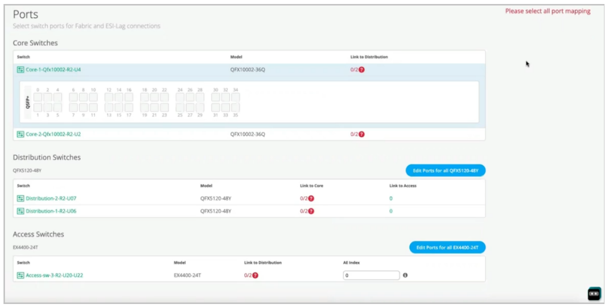

Ports

Select switch ports for Fabric and ESI-Lag connections

Please select all port mapping

Core Switches

Switch	Model	Link to Distribution
Core-1-Qfx10002-R2-U4	QFX10002-36Q	0/2
Core-2-Qfx10002-R2-U2	QFX10002-36Q	0/2

Distribution Switches

QFX5120-48Y

Edit Ports for all QFX5120-48Y

Switch	Model	Link to Core	Link to Access
Distribution-2-R2-U07	QFX5120-48Y	0/2	0
Distribution-1-R2-U06	QFX5120-48Y	0/2	0

Access Switches

EX4400-24T

Edit Ports for all EX4400-24T

Switch	Model	Link to Distribution	AE Index
Access-sw-3-R2-U20-U22	EX4400-24T	0/2	0

11. Configure the switch ports in the core layer as described below:

    a. Select a switch in the Core section to open the switch port panel.

    b. From the port panel of the core switch, select a port that you want to configure.

    c. Specify a port type (for example, ge or xe).

    d. Choose the distribution switch on which the link should terminate. You need to configure all the ports that need to be part of the campus fabric.

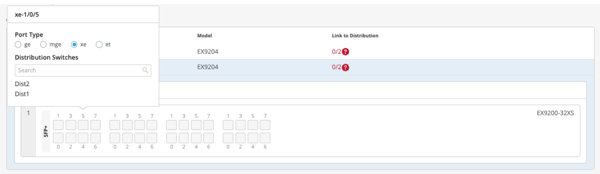

xe-1/0/5

Port Type

○ ge ○ mge ● xe ○ et

Distribution Switches

Search

Dist2

Dist1

Model	Link to Distribution
EX9204	0/2
EX9204	0/2

To configure switch ports in the distribution layer:

a. Select a switch in the Distribution section to open the switch port panel.

b. From the port panel of the switch, select a port that you want to configure.

c. Specify a port type (for example, ge or xe).

d. Select:

- **Link to Core** to connect the port to a core switch.
- **Link to Access** to connect the port to an access switch.

e. Select the core or access switch (based on the selection in the previous step) on which the link should terminate. You need to configure all the ports that need to be part of the campus fabric.

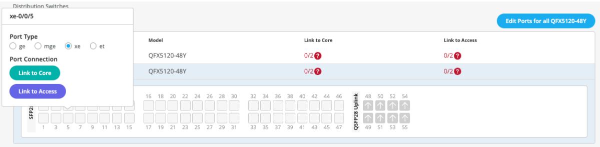

Distribution Switches

xe-0/0/5

Port Type

○ ge ○ mge ● xe ○ et

Port Connection

Link to Core

Link to Access

Edit Ports for all QFX5120-48Y

Model	Link to Core	Link to Access
QFX5120-48Y	0/2	0/2
QFX5120-48Y	0/2	0/2

To configure the switch ports in the access layer:

a. Select a switch in the Access section to open the switch port panel.

b. From the port panel of the switch, select a port that you want to configure.

c. Specify a port type (for example, `ge` or `xe`).

   In case the access layer uses a Virtual Chassis (VC), you can configure ports on the Primary and Backup tabs.

For the access switches, select only those interfaces that should be used to interconnect with the distribution switch. The system bundles all interfaces into a single Ethernet bundle through the AE index option. You can specify an AE index value for the access devices.

If you want to view the configuration and status information of a specific port, hover over the numbered box representing that port in the port panel UI.

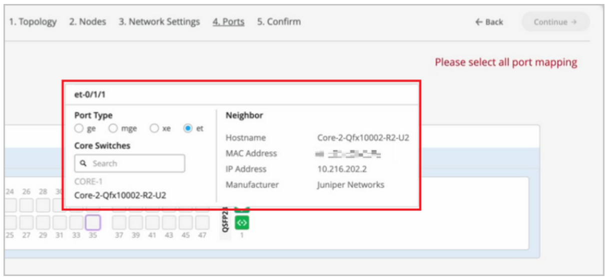

Please select all port mapping

et-0/1/1

Port Type
○ ge ○ mge ○ xe ● et

Core Switches
Search
CORE-1
Core-2-Qfx10002-R2-U2

Neighbor

Hostname	Core-2-Qfx10002-R2-U2
MAC Address	[redacted]
IP Address	10.216.202.2
Manufacturer	Juniper Networks

12. Click Continue to go to the **Confirmation** tab.
13. Click each switch icon to view and verify the configuration.
14. After verifying the configuration, click **Apply Changes** > **Confirm**.
    This step saves the campus fabric configuration to the Mist cloud and applies it to the switches. If the switches are offline, the configuration will be applied to them when they come online next time. A switch might take up to 10 minutes to complete the configuration.
15. Click **Close Campus Fabric Configuration**.

    After Mist builds the campus fabric, or while it is building the fabric, you can download the connection table. The connection table represents the physical layout of the campus fabric. You can use this table to validate all switch interconnects for the devices participating in the physical campus fabric build. Click **Connection Table** to download it (.csv format).

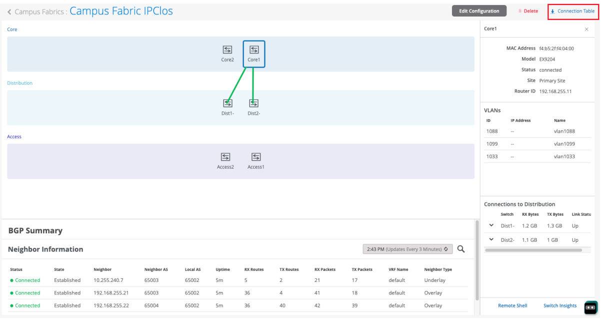

BGP Summary

Neighbor Information

2:43 PM (Updates Every 3 Minutes)

Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
● Connected	Established	10.255.240.7	65003	65002	5m	5	2	21	17	default	Underlay
● Connected	Established	192.168.255.21	65003	65002	5m	36	4	41	18	default	Overlay
● Connected	Established	192.168.255.22	65004	65002	5m	36	40	42	39	default	Overlay

Core1

MAC Address	f4:b5:2f:f4:04:00
Model	EX9204
Status	connected
Site	Primary Site
Router ID	192.168.255.11

VLANs

ID	IP Address	Name
1088	--	vlan1088
1099	--	vlan1099
1033	--	vlan1033

Connections to Distribution

Switch	RX Bytes	TX Bytes	Link Statu
Dist1-	1.2 GB	1.3 GB	Up
Dist2-	1.1 GB	1 GB	Up

16. Verify the campus fabric configuration. To verify, follow the steps listed in the **Verification** section of Campus Fabric Core Distribution CRB (JVD) and Campus Fabric Core-Distribution ERB (JVD).

    For a demo, watch the following video:

    Video: Deployment of Campus Fabric Core Distribution

Configure Campus Fabric IP Clos

SUMMARY

Follow these steps to set up a campus fabric IP Clos architecture, which enables Juniper Mist™ to provide integrated routing and bridging interfaces.

Juniper Networks campus fabrics provide a single, standards-based EVPN-VXLAN solution that you can deploy on any campus.

The campus fabric IP Clos architecture pushes VXLAN L2 gateway functionality to the access layer. This model is also called end-to-end, given that VXLAN tunnels terminate at the access layer.

The campus fabric IP Clos architecture supports Group Based Policies (GBPs) that enable you to achieve micro segmentation in the network. The GBP option gives you a practical way to create network access policies that are independent of the underlying network topology. In a GBP, you match a user group tag to a resource group tag to provide the specified users access to the specified resources. See "Create a Switch Configuration Template" on page 22 to learn how to configure GBP on switches.

In a campus fabric IP Clos architecture, Mist provisions layer 3 (L3) integrated routing and bridging (IRB) interfaces on the access layer. All the access switches are configured with the same IP address for each L3 subnet. The end users terminating on the access layer have the default gateway set to the IRB address shared by all access layer devices. This deployment model utilizes anycast addressing for all devices participating in the L3 subnet. This deployment model provides a smaller blast radius for broadcast traffic and is ideal for east-west traffic patterns and IP Multicast environments.

For more detailed information about IP Clos architecture and its deployment, see Campus Fabric IP Clos Using Mist Wired Assurance—Juniper Validated Design (JVD).

NOTE: In topologies that are built in Mist cloud after the May 2025 updates, Mist automatically detects and reports any EVPN loops and duplicate MAC addresses. These issues are displayed on the switch Insights page.

- **EVPN loop detection**—EVPN-VXLAN lightweight PE-CE loop detection helps in detecting and breaking LAN Ethernet loops on downstream leaf-to-server or access ports. This feature can detect loops caused by issues such as miswired fabric components or third-party switches incorrectly connected to the fabric.For this feature to work, the switch must run the Junos OS version 24.4R1 or later. For more information, refer to EVPN-VXLAN Lightweight Leaf to Server Loop Detection.

- **Duplicate MAC address detection**—Identifies and mitigates issues arising from MAC address movement (MAC mobility) between different interfaces or devices in EVPN environments. While some MAC mobility is expected (for example, when a device actually moves), rapid changes might indicate issues such as network loops or misconfigurations. For more information, refer to Configuring Loop Detection for Duplicate MAC Addresses.

Campus Fabric Configuration Best Practices

- Configure VLANs at switch template level and import them while configuring campus fabric. Template must be the single source of truth for all VLANs and port profiles unless specifically required at the switch or site level.

- At the access layer, avoid using trunk port profiles that allow all VLANs, unless explicitly required.

- Create VRF and VRF network configuration via the campus fabric, not via switch templates.

- Create port assignments per role and overwrite the configuration on individual device as needed.

- Manage the DHCP relay configuration via the campus fabric workflow, except for the service block devices.

To configure campus fabric IP Clos:

1. Click **Organization** > **Campus Fabric.**
2. If you want to create the campus fabric for a site, select the site from the drop-down list beside the page heading. If you want to create the campus fabric for the entire organization, select **Entire Org** from the drop-down list.

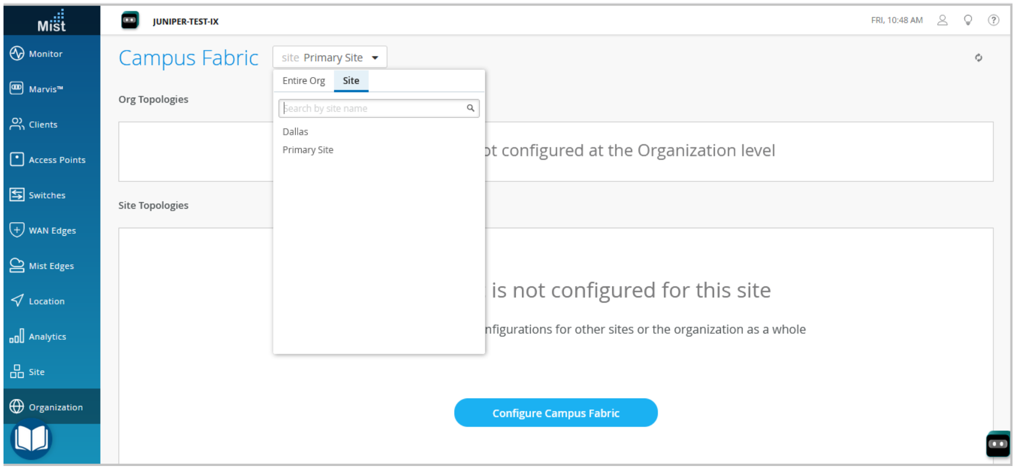

ot configured at the Organization level

is not configured for this site

nfigurations for other sites or the organization as a whole

You can use an organization-level campus fabric topology to build a campus-wide architecture with multiple buildings. Otherwise, build a site-specific campus fabric with a single set of core, distribution, and access switches.

3. Click whichever option is relevant. Click the:

   - **Configure Campus Fabric** button (displayed if the site doesn't have a campus fabric configuration associated with it).

   - **Create Campus Fabric** button (displayed if the site already has at least one campus fabric configuration associated with it).

   The **Topology** tab is displayed.

4. Select the topology type **Campus Fabric IP Clos**.

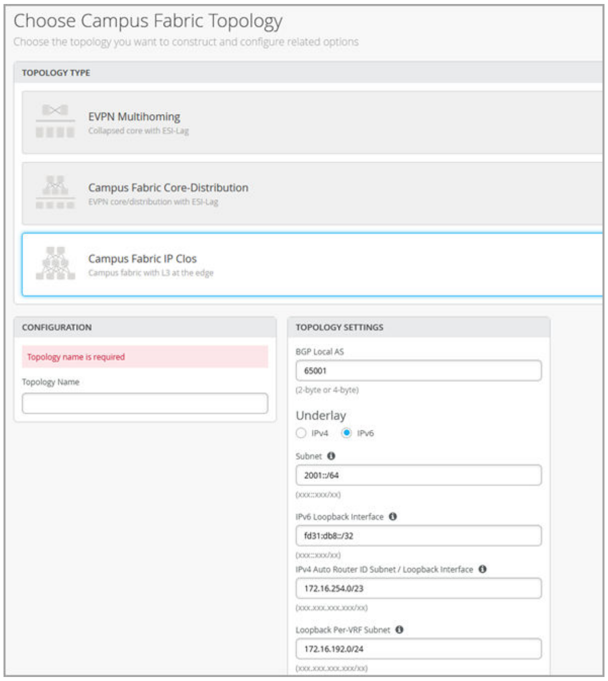
Choose Campus Fabric Topology
Choose the topology you want to construct and configure related options
TOPOLOGY TYPE
EVPN Multihoming
Collapsed core with ESI-Lag
Campus Fabric Core-Distribution
EVPN core/distribution with ESI-Lag
Campus Fabric IP Clos
Campus fabric with L3 at the edge
CONFIGURATION
TOPOLOGY SETTINGS

Topology name is required

Topology Name

BGP Local AS: 65001 (2-byte or 4-byte)

Underlay: ○ IPv4 ● IPv6

Subnet: 2001::/64 (xxx::xxx/xx)

IPv6 Loopback Interface: fd31:db8::/32 (xxx::xxx/xx)

IPv4 Auto Router ID Subnet / Loopback Interface: 172.16.254.0/23 (xxx.xxx.xxx.xxx/xx)

Loopback Per-VRF Subnet: 172.16.192.0/24 (xxx.xxx.xxx.xxx/xx)

5. Configure the topology name and other settings on the **Topology** tab, as described below:

> **NOTE**: We recommend that you use the default settings on this screen unless they conflict with any networks attached to the campus fabric. The point-to-point links between each layer utilize /31 addressing to conserve addresses.

a. In the **CONFIGURATION** section, enter the following:

   - **Topology Name**—Enter a name for the topology.

b. **(If you don't want to use the default settings)** In the **TOPOLOGY SETTINGS** section, enter the following:

   - **BGP Local AS**—Represents the starting point of private BGP AS numbers that are automatically allocated to each device. You can use any private BGP AS number range that suits your deployment. Mist provisions the routing policy so that only the loopback IP addresses are exchanged in the underlay of the fabric.

   - **Underlay**—Select an internet protocol version for the underlay. Options are IPv4 and IPv6.

   - **Subnet**— The range of IP addresses that Mist uses for point-to-point links between devices. You can use a range that suits your deployment. Mist breaks this subnet into /31 subnet addressing per link. You can modify this number to suit the specific deployment scale. For example, a /24 network would provide up to 128 point-to-point /31 subnets.

   - **IPv6 Loopback Interface**—Specify an IPv6 loopback interface subnet, which is used to autoconfigure IPv6 loopback interface on each device in the fabric.

   - **IPv4 Auto Router ID Subnet / Loopback Interface**—Mist uses this subnet to automatically assign a router ID to each device in the fabric (including access devices irrespective of whether they are configured with EVPN or not). Router IDs are loopback interfaces (lo0.0) used for overlay peering between devices. For new topologies, this field auto-populates a default subnet value (172.16.254.0/23), which you can modify. When you edit an existing topology, this field doesn't populate a default value. The router ID is used as an identifier when deploying routing protocols such as BGP.

     You can overwrite the automatically assigned router ID by manually configuring a loopback interface in the Router ID field on the Routing tile on the switch configuration page (**Switches** > ***Switch Name***). However, if you modify the campus fabric configuration afterwards, Mist performs the automatic assignment of the router ID again, replacing the manually configured loopback interface.

   - **Loopback per-VRF subnet**—Mist uses this subnet to automatically configure loopback interfaces (lo0.x) per virtual routing and forwarding (VRF) instance that is used for services such as DHCP relay. For new topologies, this field auto-populates a default subnet value (172.16.192.0/24), which you can modify. This field supports a /19 or smaller subnet (for example, /24). When you edit an existing topology, this field doesn't populate a default value.

6. Click **Continue** to go to the **Nodes** tab, where you can select devices that form part of the campus fabric IP Clos deployment.

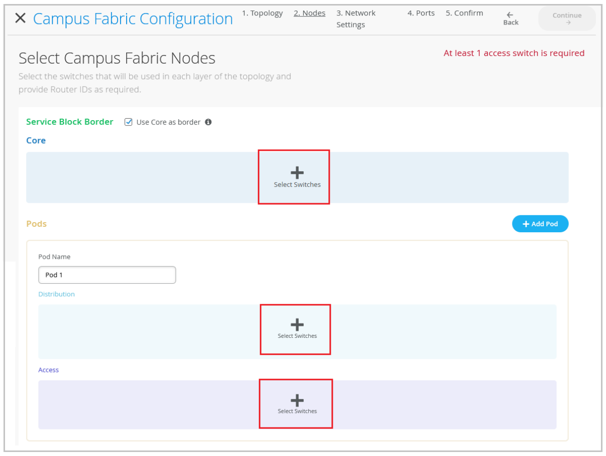

Select Campus Fabric Nodes

Select the switches that will be used in each layer of the topology and provide Router IDs as required.

At least 1 access switch is required

☑ Use Core as border

Pod Name: Pod 1

7. Add switches to the Core, Distribution, and Access layer sections as required.

   To add the switches:

   a. Click **Select Switches** in the section to which you want to add switches.

   b. Select the switches that you want to add to the campus fabric.

   c. Click **Select**.

   We recommend that you validate the presence of each device in the switch inventory before creating the campus fabric.

   By default, Mist configures the core switches to function as border nodes that run the services block functionality. In a campus fabric topology, border nodes interconnect external devices such as firewalls, routers, or critical devices. External services or devices (for example, DHCP and RADIUS servers) connect to the campus fabric through border nodes. If you want to offload this task from the core switches and use dedicated switches as border nodes, clear the **Use Core as border** checkbox on the upper left of the page. You can then add up to two switches as dedicated border nodes. The minimum number of dedicated border nodes required is one.

   Also, Mist provides pods for improved scalability. Your access and distribution devices are grouped into pods. A pod could represent a building. For example, you can create a pod for each of the buildings in your site and create connections between the access and the distribution devices in that pod. You do not have to connect the same set of access devices to the distribution devices across multiple buildings. You can create multiple pods by clicking **+Add Nodes**.

   You need only one connection between a pod and the core switch. You do not have to connect each distribution switch in a pod to all the core switches used. In an IPClos topology, you need only one connection between each core and distribution pair and between each distribution and access pair.

8. After selecting the switches, click **Continue** to go to the **Network Settings** tab, where you can configure the networks.

9. Configure the network settings, as described below.

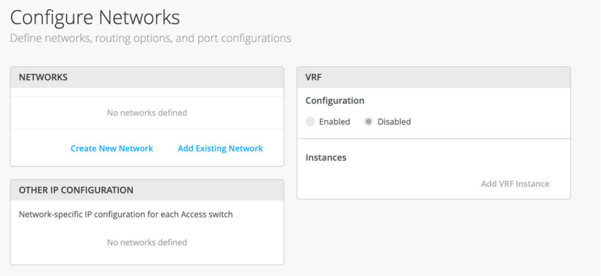

Configure Networks

Define networks, routing options, and port configurations

NETWORKS

No networks defined

Create New Network   Add Existing Network

OTHER IP CONFIGURATION

Network-specific IP configuration for each Access switch

No networks defined

VRF

Configuration

○ Enabled ● Disabled

Instances

Add VRF Instance

a. From the NETWORKS tile, add networks or VLANs to the configuration. You can either create a new network or import the network from the switch template defined in the **Organization > Switch** templates page.

   To add a new VLAN, click **Create New Network** and configure the VLANs. The settings include a name, VLAN ID, and a subnet. You can specify IPv4 or IPv6 addresses for the subnet.

   You can optionally configure IPv4 and IPv6 Anycast Gateway addresses in addition to any IPv4 or IPv6 subnets you have configured. The Mist UI uses those gateways as the Anycast IP address assignment across all Access and Distribution switches in the Campus Fabric configuration.

   > **NOTE**: If the Anycast Gateway fields are left empty, the Mist UI uses the existing logic, which is to use the first IP address in the subnet as the Anycast address.

   To import VLANs from the template:

   i. Click **Add Existing Network**.

   ii. Select a switch template from the **Template** drop-down list to view the VLANs available in that template.

   iii. Select the required VLAN from the displayed list, and click the ✓ mark.

   VLANs are mapped to Virtual Network Identifiers (VNIs). You can optionally map the VLANs to VRF instances to logically separate the traffic.

b. Review the settings on the OTHER IP CONFIGURATION tile. This tile populates the settings automatically after you specify the networks in the NETWORKS section.

   Mist provides automatic IP addressing of IRB for each of the VLANs. Then, the port profile associates the VLAN with the specified ports.

c. Optionally, configure VRF instances. Mist recommends using VRFs in network segments where traffic isolation and overlapping IP address spaces are required. By default, Mist places all VLANs in the default VRF. The VRF option allows you to group common VLANs into the same VRF or separate VRFs depending on traffic isolation requirements. All VLANs within each VRF have full connectivity with each other and with other external networking resources. A common use case is the isolation of guest wireless traffic from most enterprise domains except Internet connectivity. By default, a campus fabric provides complete isolation between VRFs, forcing inter-VRF communications to traverse a firewall. If you require inter-VRF communication, you need to include extra routes to the VRF. The extra route could be a default route that instructs the campus fabric to use an external router. It could also be a firewall for further security inspection or routing capabilities.

To create a VRF:

i. On the VRF tile, click **Add VRF Instance** and specify the settings. The settings include a name for the VRF and the networks to be associated with the VRF.

ii. To add extra routes, click the **Add Extra Routes** link on the **New VRF Instance** page and specify the route. You can specify IPv4 or IPv6 addresses.

d. On the DHCP RELAY tile, configure the DHCP relay settings. You have the following options:

- Enabled—Configures DHCP relay on all the IRB-enabled devices in campus fabric. This option allows you to enable DHCP Relay on networks that you selected. The network will be populated inside the DHCP Relay tile as long as it is listed on the Networks tab on the same page.

- Disabled—Disable DHCP relay on the devices in campus fabric. When you select this option, the DHCP relay is disabled on all the IRB-enabled devices. You should carefully select this option as this will remove the locally defined DHCP Relay on the Switch Detail page.

- None—This option is automatically selected when the campus fabric topology has a mix of devices in terms of the DHCP relay configuration; that is, some devices have the DHCP relay enabled, some have it disabled, and some do not have it defined.

If you want to remove all locally defined DHCP Relay networks, select **Enabled** and then choose **Remove all existing device level DHCP Networks**. You can simplify your DHCP Relay deployment by centralizing any configuration change from the campus fabric workflow.

If you enable DHCP relay in a campus fabric configuration, it is enabled on all the IRB-defined devices in the fabric and disabled on the rest of the devices. For example, in Campus Fabric IP Clos edge topologies, DHCP is enabled on access devices and disabled on the rest.

10. Click **Continue** to go to the **Ports** tab, where you can configure the ports and create a connection between the core, distribution, and access layer switches.

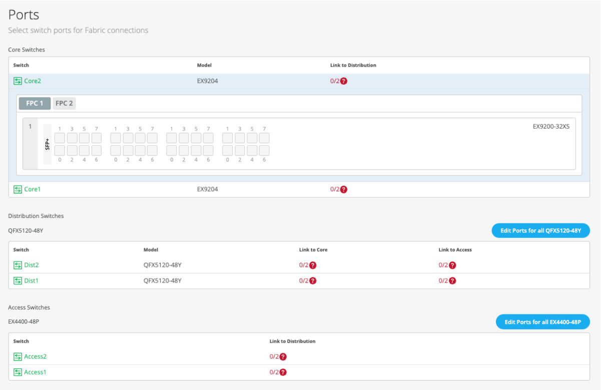

Ports

Select switch ports for Fabric connections

Core Switches

Switch	Model	Link to Distribution
Core2	EX9204	0/2 ?
Core1	EX9204	0/2 ?

Distribution Switches

QFX5120-48Y

Switch	Model	Link to Core	Link to Access
Dist2	QFX5120-48Y	0/2 ?	0/2 ?
Dist1	QFX5120-48Y	0/2 ?	0/2 ?

Access Switches

EX4400-48P

Switch	Link to Distribution
Access2	0/2 ?
Access1	0/2 ?

11. Configure the switch ports in the core layer as described below:

    a. Select a switch in the Core section to open the switch port panel.

    b. From the port panel of the core switch, select a port that you want to configure.

    c. Specify a port type (for example, ge or xe).

    d. Choose the distribution switch on which the link should terminate. You need to configure all the ports that need to be part of the campus fabric.

    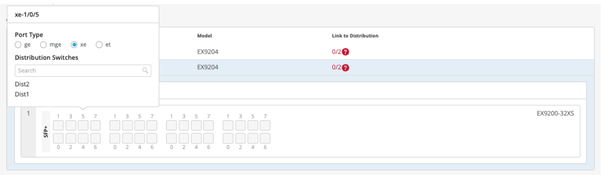

    xe-1/0/5

    Port Type

    ○ ge ○ mge ● xe ○ et

    Distribution Switches

    Search

    Dist2

    Dist1

Model	Link to Distribution
EX9204	0/2 ?
EX9204	0/2 ?

    To configure switch ports in the distribution layer:

    a. Select a switch in the Distribution section to open the switch port panel.

    b. From the port panel of the switch, select a port that you want to configure.

    c. Specify a port type (example: ge or xe).

    d. Select:

    - **Link to Core** to connect the port to a core switch.

    - **Link to Access** to connect the port to an access switch.

    e. Select the core or access switch (based on the selection in the previous step) on which the link should terminate. You need to configure all the ports that need to be part of the campus fabric.

    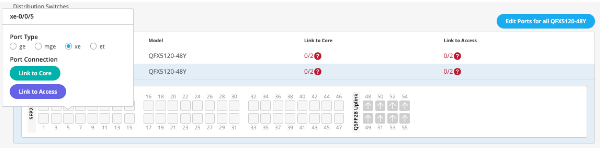

    Distribution Switches

    xe-0/0/5

    Port Type

    ○ ge ○ mge ● xe ○ et

    Port Connection

Model	Link to Core	Link to Access
QFX5120-48Y	0/2 ?	0/2 ?
QFX5120-48Y	0/2 ?	0/2 ?

To configure switch ports in the access layer:

a. Select a switch in the Access section to open the switch port panel.

b. From the port panel of the switch, select a port that you want to configure.

c. Specify a port type (example: ge and xe).

d. Choose the distribution switch on which the link should terminate. You need to configure all the ports that need to be part of the campus fabric.

If you want to view the configuration and status information of a specific port, hover over the numbered box representing that port in the port panel UI.

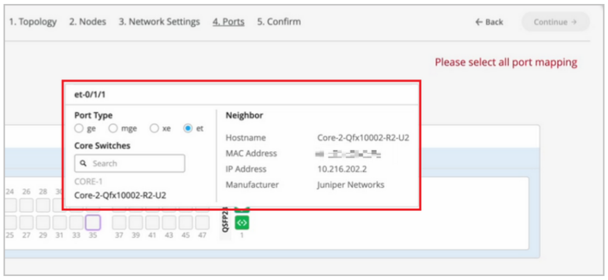

Please select all port mapping

et-0/1/1

Port Type
○ ge ○ mge ○ xe ● et

Core Switches
Search
CORE-1
Core-2-Qfx10002-R2-U2

Neighbor

Hostname	Core-2-Qfx10002-R2-U2
MAC Address	[redacted]
IP Address	10.216.202.2
Manufacturer	Juniper Networks

12. Click Continue to go to the **Confirmation** tab.

13. Click each switch icon to view and verify the configuration.

14. After verifying the configuration, click **Apply Changes** > **Confirm**.

    The campus fabric configuration is saved to the Mist cloud. The configuration is immediately applied to the switches if they are online. If the switches are offline, the configuration will be applied to them when they come online next time. A switch might take up to 10 minutes to complete the configuration.

15. Click **Close Campus Fabric Configuration**.

    Once the campus fabric is built or is in the process of being built, you can download the connection table, which represents the physical layout of the campus fabric. You can use this table to validate all switch interconnects for the devices participating in the physical campus fabric build. Click **Connection Table** to download it (.csv format).

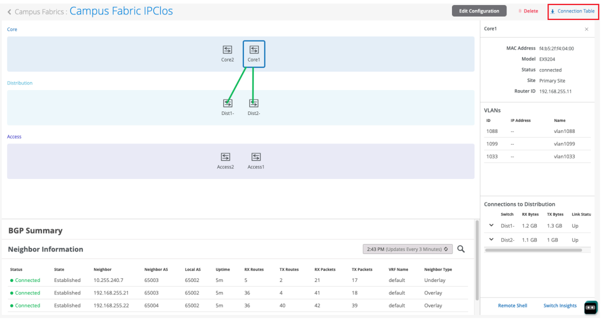

Core1

MAC Address	f4:b5:2f:f4:04:00
Model	EX9204
Status	connected
Site	Primary Site
Router ID	192.168.255.11

VLANs

ID	IP Address	Name
1088	--	vlan1088
1099	--	vlan1099
1033	--	vlan1033

Connections to Distribution

Switch	RX Bytes	TX Bytes	Link Statu
Dist1-	1.2 GB	1.3 GB	Up
Dist2-	1.1 GB	1 GB	Up

BGP Summary

Neighbor Information

2:43 PM (Updates Every 3 Minutes)

Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
● Connected	Established	10.255.240.7	65003	65002	5m	5	2	21	17	default	Underlay
● Connected	Established	192.168.255.21	65003	65002	5m	36	4	41	18	default	Overlay
● Connected	Established	192.168.255.22	65004	65002	5m	36	40	42	39	default	Overlay

16. Verify the campus fabric configuration. To verify, follow the steps listed in the **Verification** section of Campus Fabric IP Clos Wired Assurance.

    For a demo of the configuration steps, watch the following video:

Video: Deployment of Campus Fabric IP Clos

Mist does not automatically peer a campus fabric with the external network. To enable external connectivity, you must manually configure the necessary settings (such as the interface and VLAN you need to peer with) on each core switch. You can configure these settings in the Additional IP Configuration section on the IP Configuration tile of the switch details page. Additionally, if the external network uses an overlay, it must be added to a VRF on the specific device and then referenced in the BGP or OSPF configuration.

After building an IP Clos campus fabric, you can integrate it with a third party gateway (such as a router or firewall) by using BGP groups. Watch the following video for more information:

Video: Integrate IP Clos Fabric Using BGP

See also: "Configure BGP on Switches via Mist" on page 89.

Manage Campus Fabric Configuration

SUMMARY

Use the instructions in this topic to view, modify, or delete an organization-level or site-level campus fabric configuration.

IN THIS SECTION

View Campus Fabric Configuration

To view a campus fabric configuration:

1. Click **Organization** > **Campus Fabric**.
2. Click the campus fabric configuration to open it.

NOTE: Site-level campus fabric configurations are displayed only if you choose the site from the scope selector placed near the page title.

3. Click **View Configuration**.
4. Navigate through the configuration tabs to view the configuration.

Modify Campus Fabric Configuration

To modify a campus fabric configuration:

1. Click **Organization** > **Campus Fabric**.
2. Click the campus fabric configuration to open it.

NOTE: Site-level campus fabric configurations are displayed only if you choose the site from the scope selector placed near the page title.

3. Click **Edit Configuration**.
4. Navigate through the configuration tabs to make the required changes.

   Refer to the following topics for more information on the configuration fields:

   - "Configure Campus Fabric EVPN Multihoming" on page 257
   - "Configure Campus Fabric IP Clos" on page 275
   - "Configure Campus Fabric IP Clos" on page 275
5. After making the changes, click **Apply Changes** on the Confirm tab.
6. Click **Confirm**.

Delete a Campus Fabric Configuration

To delete a campus fabric configuration from Mist:

1. Click **Organization** > **Campus Fabric**.
2. Click the campus fabric configuration to open it.

NOTE: Site-level campus fabric configurations are displayed only if you choose the site from the scope selector placed near the page title.

3. Click **Delete**.
   A confirmation window (Confirm Delete) appears.
4. In the Topology Name field on the Confirm Delete window, specify the name of the campus fabric configuration to be deleted.
5. Click **Delete**.

Group-Based Policies

SUMMARY

Learn about Group-Based Policies, create GBP tags, and set up switch policies to allow or deny access to tagged packets.

IN THIS SECTION

Group-Based Policies Overview

A group-based policy (GBP) helps you achieve microsegmentation and macrosegmentation to secure data and assets in Virtual extensible Local Area Network (VXLAN) architecture. GBP leverages the underlying VXLAN technology to provide location-agnostic endpoint access control. GBP allows you to implement consistent security policies across the enterprise network domains, and simplifies your network configuration as it spares you the need to configure large number of firewall filters on all your switches. GBP blocks lateral threats by ensuring consistent application of security group policies throughout the network, regardless of the location of endpoints or users.

GBPs can be utilized in Juniper Mist campus fabric IP Clos topologies to achieve micro segmentation in the network. The GBP option gives you a practical way to create network access policies that are independent of the underlying network topology.

VXLAN-GBP works by leveraging reserved fields in the VXLAN header for use as a Scalable Group Tag (SGT). You can use the SGTs to match conditions in firewall filter rules. Using an SGT is more robust than using port or Media Access Control (MAC) addresses to achieve comparable results. SGTs can be assigned statically (by configuring the switch on a per port or per MAC basis), or they can be configured on the Remote Authentication Dial in User Service (RADIUS) server and pushed to the switch through 802.1X when the user is authenticated.

The segmentation enabled by VXLAN-GBP is especially useful in campus VXLAN environments because it provides a practical way to create network access policies that are independent of the underlying network topology. Segmentation simplifies the design and implementation phases of developing network-application and endpoint-device security policies.

Watch the following video for a quick overview of GBP:

Video: Campus Fabric GBP Microsegmentation

On the Mist portal, you can configure GBP using the switch templates (**Organization** > **Switch Templates**), or directly from the switch configuration page (**Switches** > ***Switch Name***). The GBP configuration involves creating GBP tags and including them in switch policies. The GBP tags enable you to group users and resources. In a GBP, you match a user group tag to a resource group tag to provide the specified users access to the specified resources.

The following video takes you through the steps involved in configuring a GBP:

Video: Group Based Policy Demo

See also: Microsegmentation with GBP Using Mist Wired Assurance.

Create Switch Policies Using GBP Tags

In the Juniper Mist portal, you'll add GBP tags that you can then reference in your switch policies to allow or block tagged frames.

NOTE: Supported platforms include EX4400, EX4100, EX4650, QFX5120-32C, and QFX5120-48Y running Junos OS Release 22.4R1 and later.

The following examples show the tags that a user created, followed by the policies that reference them.

Figure 20: GBP Tags

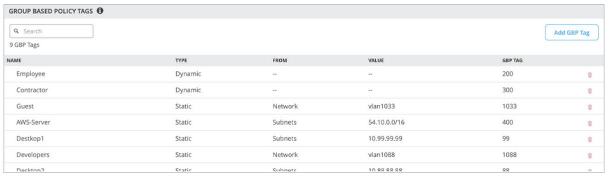

9 GBP Tags

NAME	TYPE	FROM	VALUE	GBP TAG
Employee	Dynamic	--	--	200
Contractor	Dynamic	--	--	300
Guest	Static	Network	vlan1033	1033
AWS-Server	Static	Subnets	54.10.0.0/16	400
Destkop1	Static	Subnets	10.99.99.99	99
Developers	Static	Network	vlan1088	1088
Desktop2	Static	Subnets	10.88.88.88	88

Figure 21: Switch Policies Using GBP Tags

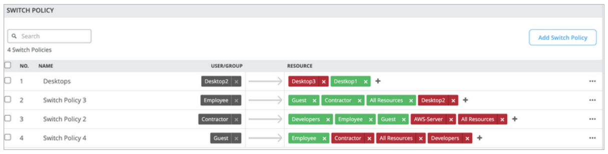

SWITCH POLICY

Search

4 Switch Policies

□	NO.	NAME	USER/GROUP		RESOURCE	
□	1	Desktops	Desktop2 ×	→	Desktop3 × Destkop1 × +	···
□	2	Switch Policy 3	Employee ×	→	Guest × Contractor × All Resources × Desktop2 × +	···
□	3	Switch Policy 2	Contractor ×	→	Developers × Employee × Guest × AWS-Server × All Resources × +	···
□	4	Switch Policy 4	Guest ×	→	Employee × Contractor × All Resources × Developers × +	···

1. Navigate to your switch template or site-level switch configuration.

   - To find a switch template—From the left menu, select **Organization** > **Wired** > **Switch Templates**. Then click the template.

   - To find a site-level switch configuration—From the left menu, select **Site** > **Wired** > **Switch Configuration**. Then click the site.

2. Scroll down to the **Group Based Policy Tags** section, and add a GBP tag:

   a. Click **Add GBP tag**.

   b. Enter the details:

      - **Name**—Enter a short, specific name, which will help you to identify this tag when creating your switch policies.

      - **Dynamic or Static**—By default, Juniper Mist chooses the Dynamic option. **If** you select Static, specify a GBP tag source by using the drop-down list.

      - **GBP Tag**—Enter value or GBP source tag for host-originated packets (range: 1 through 65535).

   c. Click **Add**.

      When you enable a tag, you'll see on-screen alert about the impact on standalone switches and virtual chassis. Read the on-screen information before proceeding.

> **NOTE**: If you configure 802.1X authentication with multiple-supplicant mode, the GBP tagging is MAC-based. If you configure 802.1X authentication with single-supplicant mode, the GBP tagging is port-based.

3. Scroll down to the **Switch Policy** section and add a policy that uses your new tag:

   a. If you're working on a site-level switch configuration, select **Override Template Defined**.

   b. Click **Add Switch Policy**.
      A new policy appears at the top of the policy list, with a default name such as Switch Policy 1.

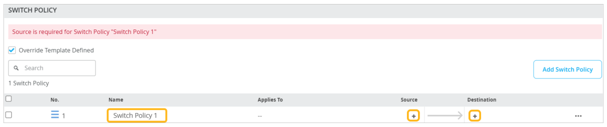

SWITCH POLICY

Source is required for Switch Policy "Switch Policy 1"

☑ Override Template Defined

Search

1 Switch Policy

□	No.	Name	Applies To	Source		Destination	
□	☰ 1	Switch Policy 1	--	+	→	+	···

   c. Click the default policy name, and then enter a short, specific name to identify this policy.

d. Under **Source**, click the **+** button, and then select a GBP tag from the list.

e. Under **Destination**, click the **+** button, and then select a GBP tag.
The destination label appears with a green background, indicating that this policy allows access to this destination.

f. If you want this policy to block access to the specified destination, click the destination button, and then click **Deny**.

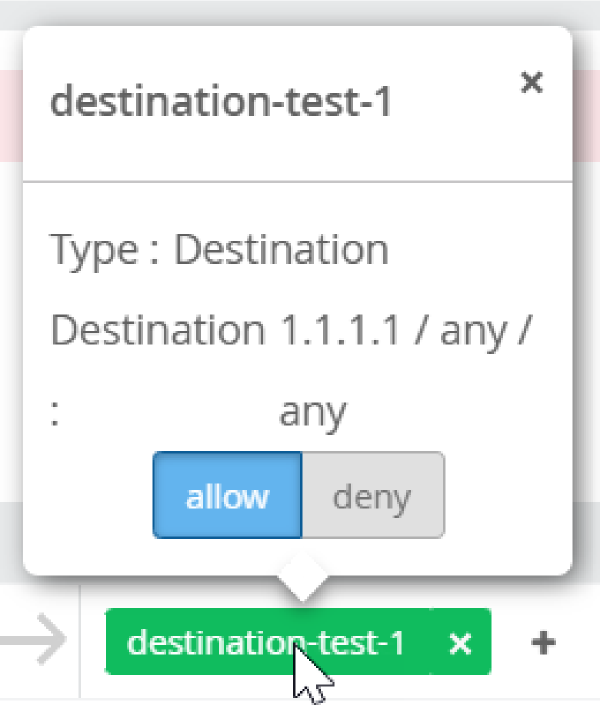

destination-test-1

Type : Destination

Destination : 1.1.1.1 / any / any

When you change the policy to deny, the destination button turns red.

> **NOTE**: The default action is allow (green button).

4. Click **Save** at the top-right corner of the configuration page.
5. Review the on-screen confirmation information, and then click **Save** at the bottom of the confirmation window.

5

CHAPTER

Switch Dashboards

IN THIS CHAPTER

Switches (List View)

Click **Switches** on the left navigation menu on the Mist portal to navigate to the Switches page.

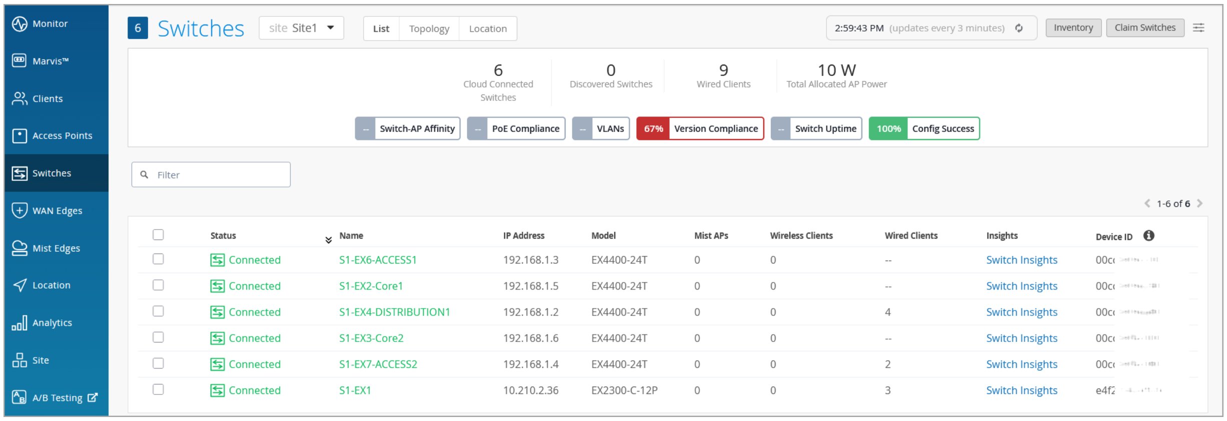

6 Cloud Connected Switches
0 Discovered Switches
9 Wired Clients
10 W Total Allocated AP Power

-- Switch-AP Affinity
-- PoE Compliance
-- VLANs
67% Version Compliance
-- Switch Uptime
100% Config Success

☐	Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	Insights	Device ID
☐	Connected	S1-EX6-ACCESS1	192.168.1.3	EX4400-24T	0	0	--	Switch Insights	00c[illegible]
☐	Connected	S1-EX2-Core1	192.168.1.5	EX4400-24T	0	0	--	Switch Insights	00c[illegible]
☐	Connected	S1-EX4-DISTRIBUTION1	192.168.1.2	EX4400-24T	0	0	4	Switch Insights	00c[illegible]
☐	Connected	S1-EX3-Core2	192.168.1.6	EX4400-24T	0	0	--	Switch Insights	00c[illegible]
☐	Connected	S1-EX7-ACCESS2	192.168.1.4	EX4400-24T	0	0	2	Switch Insights	00c[illegible]
☐	Connected	S1-EX1	10.210.2.36	EX2300-C-12P	0	0	3	Switch Insights	e4f2[illegible]

This page lists the switches that have been onboarded to Mist and assigned to a specific site.

The list view includes the key switch details such as connection status, switch name, IP address, MAC address, switch model, connected APs, wired and wireless clients, device ID, and loopback IP address.

NOTE: The switches that are not assigned to a site are not listed here. To view all switches claimed to your organization, including those not yet assigned to a site, you must navigate to the inventory page (**Organization** > **Inventory** > **Switches**).

At the top of the Switches page, you can find "Switch Metrics" on page 293, which help you track the switch performance against certain compliance parameters.

You can customize the switch list view by adding or removing columns. To do that, click the Table Settings button on the upper right of the page, select or clear the fields based on your requirement, and then click **Close**.

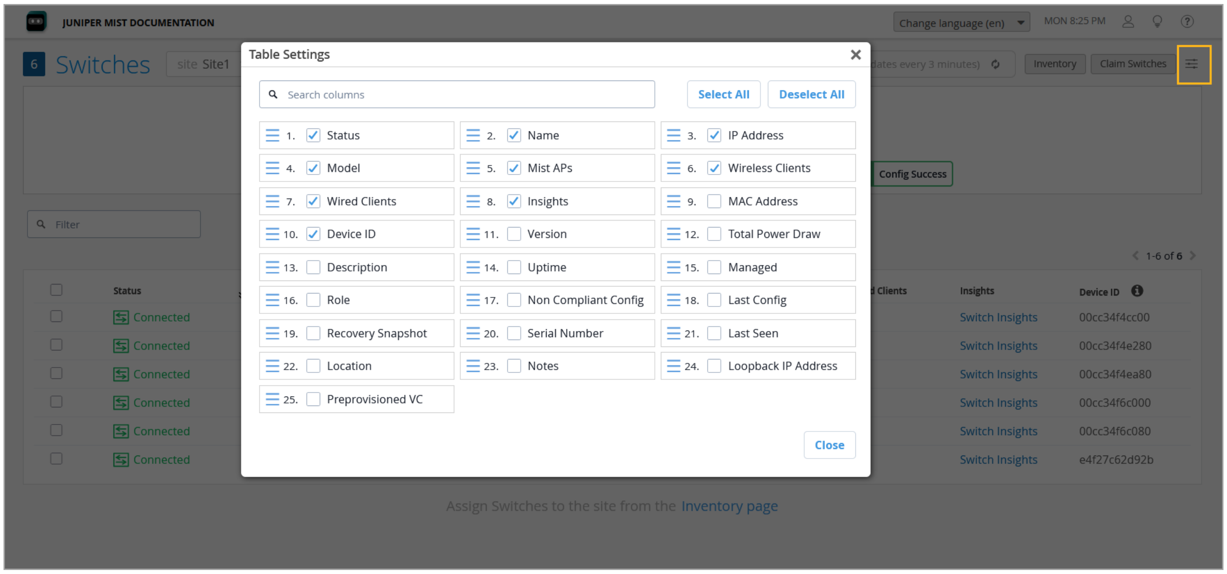

Table Settings

Search columns

Select All | Deselect All

1. ☑ Status
2. ☑ Name
3. ☑ IP Address
4. ☑ Model
5. ☑ Mist APs
6. ☑ Wireless Clients
7. ☑ Wired Clients
8. ☑ Insights
9. ☐ MAC Address
10. ☑ Device ID
11. ☐ Version
12. ☐ Total Power Draw
13. ☐ Description
14. ☐ Uptime
15. ☐ Managed
16. ☐ Role
17. ☐ Non Compliant Config
18. ☐ Last Config
19. ☐ Recovery Snapshot
20. ☐ Serial Number
21. ☐ Last Seen
22. ☐ Location
23. ☐ Notes
24. ☐ Loopback IP Address
25. ☐ Preprovisioned VC

Close

You can also access Switch Insights from this page.

In addition, you can perform multiple operations on switches, including the following:

- Upgrade Firmware—Upgrade the Junos OS software on switches. Fore more information, refer to "Upgrade Junos OS on Switches" on page 153.
- Bulk Upload Configuration—Configure switch settings by uploading them through a CSV file. Fore more information, refer to "Configure Switch-Specific Settings Using the Bulk Upload Option" on page 25.
- Assign Switch Role—Assign a role to a switch so that it can be managed by role-based configuration rules in your configuration template. Fore more information, refer to "Assign a Role to Switches" on page 174.
- Snapshot Device—Create a recovery snapshot for switches. A recovery snapshot stored in OAM (Operations, Administration, and Maintenance) volume holds a full backup that can be used in case something goes wrong with a Junos volume. Fore more information, refer to "Create Recovery Snapshot for a Switch" on page 172.
- Disable Switch Configuration—When you disable a switch configuration, the switch can no longer be managed through Mist.
- Assign to Site—Assign switches to a site.
- Rename—Rename switches.
- Release—For more information, refer to "Release a Switch from Inventory" on page 194.
- Upload Images—Upload switch photos. These photos get associated with the switch and can be managed from the PROPERTIES section on the switch details page.

Switch Metrics

SUMMARY

Get familiar with the metrics that you can use to monitor your switches and address issues with excess APs, over-allocated PoE, low uptime, and more.

The metrics on the Switches page help you track the switch performance against certain compliance parameters.

To view switch metrics, click **Switches** on the left navigation pane on the Mist portal. Metrics for which there is no data appear grey (usually, because the feature has not been configured).

Figure 22: Switch Metrics on the Switches Page

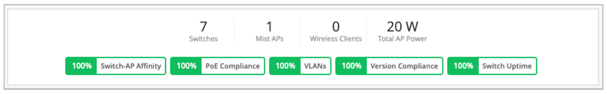

Here's a list of the switch metrics you can track:

- **Switch-AP Affinity**—This indicator shows the weighted percentage of the switches for which the number of APs connected exceeds the threshold configured. By default, the Switch-AP Affinity threshold is set to 12 APs per switch. You can configure a threshold value for the number of APs per switch to be considered in the Switch-AP Affinity metric calculation. To configure the AP threshold, click the Switch-AP Affinity indicator, and then click the hamburger icon on the right of the Switch-AP Affinity section.

- **PoE Compliance**—This indicator shows the percentage of APs that have the required 802.3at power. PoE compliance is impacted when the APs draw more power from the switch than what is allocated.

- **VLANs**—This indicator shows the percentage of APs for which all the wired VLANs are active. Click this indicator to view the list of switches and APs that have inactive or missing VLANs, along with the port information.

- **Version Compliance**— This indicator shows the percentage of switches that have the same Junos software version (per switch model). To achieve 100 percent version compliance, you must ensure that all the switches in your site run the same Junos version per switch model.

- **Switch Uptime**—This indicator shows the percentage of time a switch was up during the past seven days, averaged across all switches. Click this indicator to view the list of switches that had less than 100 percent uptime during the past 7 days.

You can click each of the switch metric indicators to get a filtered view and quickly access each device dashboard.

For more details on the switch metrics, watch the following video:

Video: Switch Metrics Overview

Switch Details

SUMMARY

Get familiar with the information and troubleshooting tools that you can find on the Switch Details page.

IN THIS SECTION

The switch details page is your ultimate go-to place on Mist for everything you need to know about a switch. You can view the status of each port and the statistics of the devices connected to the switch, access the switch configuration, review and modify the configuration, and track how the switch is performing against key metrics that matter to you.

The switch details page also has the tools that help you with switch testing and troubleshooting.

Finding the Switch Details Page

From the left menu, select **Switches**. Then click a switch in the list.

Here's an example of the details page of a switch named **ld-cup-idf-a-sw22**.

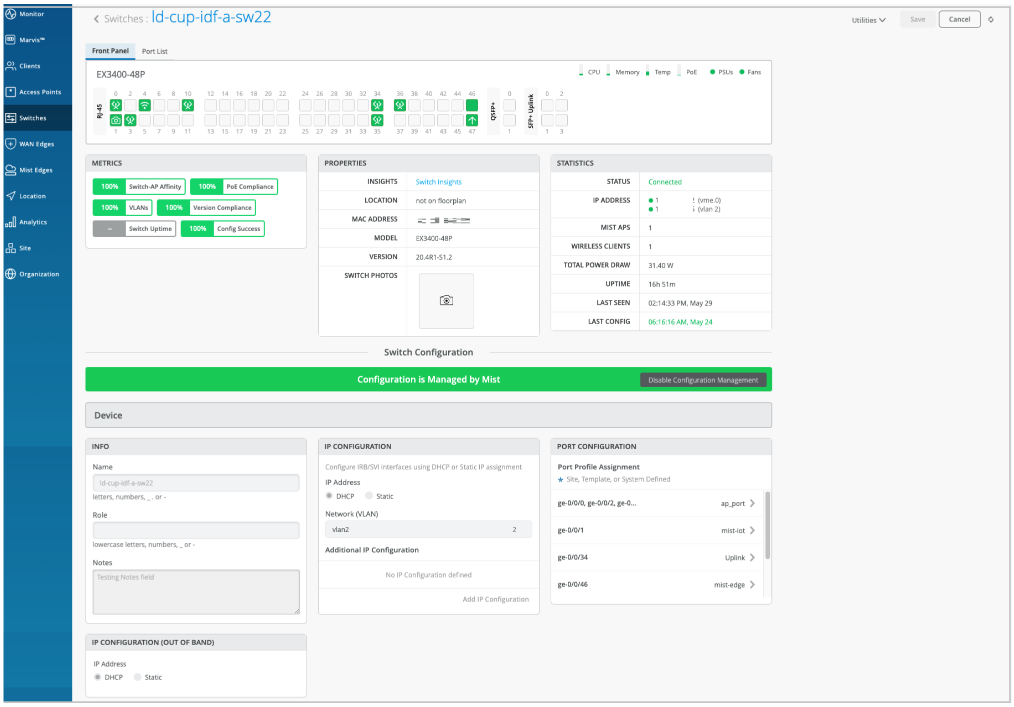

Switches : ld-cup-idf-a-sw22

EX3400-48P

METRICS

100% Switch-AP Affinity
100% PoE Compliance
100% VLANs
100% Version Compliance
-- Switch Uptime
100% Config Success

PROPERTIES

INSIGHTS	Switch Insights
LOCATION	not on floorplan
MAC ADDRESS	[redacted]
MODEL	EX3400-48P
VERSION	20.4R1-S1.2
SWITCH PHOTOS	

STATISTICS

STATUS	Connected
IP ADDRESS	● 1[redacted] (vme.0) ● 1[redacted] (vlan 2)
MIST APS	1
WIRELESS CLIENTS	1
TOTAL POWER DRAW	31.40 W
UPTIME	16h 51m
LAST SEEN	02:14:33 PM, May 29
LAST CONFIG	06:16:16 AM, May 24

Switch Configuration

Configuration is Managed by Mist

Device

INFO

Name: ld-cup-idf-a-sw22
letters, numbers, _ . or -
Role:
lowercase letters, numbers, _ or -
Notes: Testing Notes field

IP CONFIGURATION

Configure IRB/SVI interfaces using DHCP or Static IP assignment
IP Address
● DHCP ○ Static
Network (VLAN): vlan2 2
Additional IP Configuration
No IP Configuration defined

PORT CONFIGURATION

Port Profile Assignment
★ Site, Template, or System Defined

ge-0/0/0, ge-0/0/2, ge-0...	ap_port
ge-0/0/1	mist-iot
ge-0/0/34	Uplink
ge-0/0/46	mist-edge

IP CONFIGURATION (OUT OF BAND)

IP Address
● DHCP ○ Static

Front Panel

When you open a switch details page, you'll find yourself on the Front Panel tab. This tab gives you a comprehensive overview of the switch's port panel.

In this Front Panel view, you get a logical representation of the switch's ports. You can view the port status, port configuration, and the clients or APs connected to each port. The following image represents a sample Front Panel.

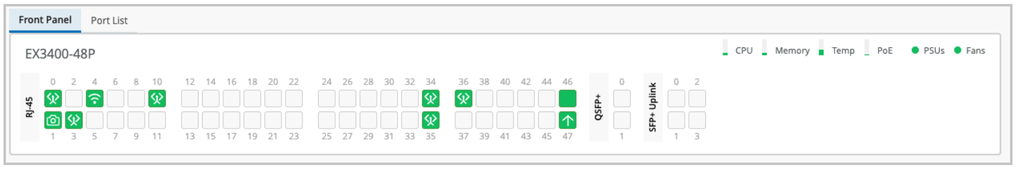

The port icons on the Front Panel view help you quickly identify the client devices or APs connected to each port and their status. The following table lists the key port icons and their descriptions:

Table 22: Port Icons and Their Descriptions

Port Icon	Description
□	The port is empty. No device is connected.
<•>	A wired client is connected.
<⚲>	A wired client is connected (trunk port).
📶	A Mist AP is connected.
📷	A camera is connected. This icon applies only to Verkada cameras.
⤢	Virtual Chassis port (VCP). A member device is connected.
🟩	The port is up. Sometimes, when a switch port is learning multiple MAC addresses on the same interface, the switch cannot identify which device is connected to the port. When that happens, the Mist portal might not display the connected device as a wired client, even though the port icon stays solid green. However, if the connected device has LLDP enabled, the portal identifies which device is connected to the port.
▢	The port is empty, with active alerts.
<•>	A wired client is connected, but the port has active alerts.

Table 22: Port Icons and Their Descriptions *(Continued)*

Port Icon	Description
📶	A Mist AP is connected, but the port has active alerts.
⬆	An uplink is connected. But the port has active alerts.

For a sneak peek into what's happening on a specific port, simply hover your mouse over the port. You can view the type of device connected to the port, the port speed, power settings, IP address, authentication status, authentication method, and much more.

Here's an example of what you might see if you hover over port ge-0/0/32:

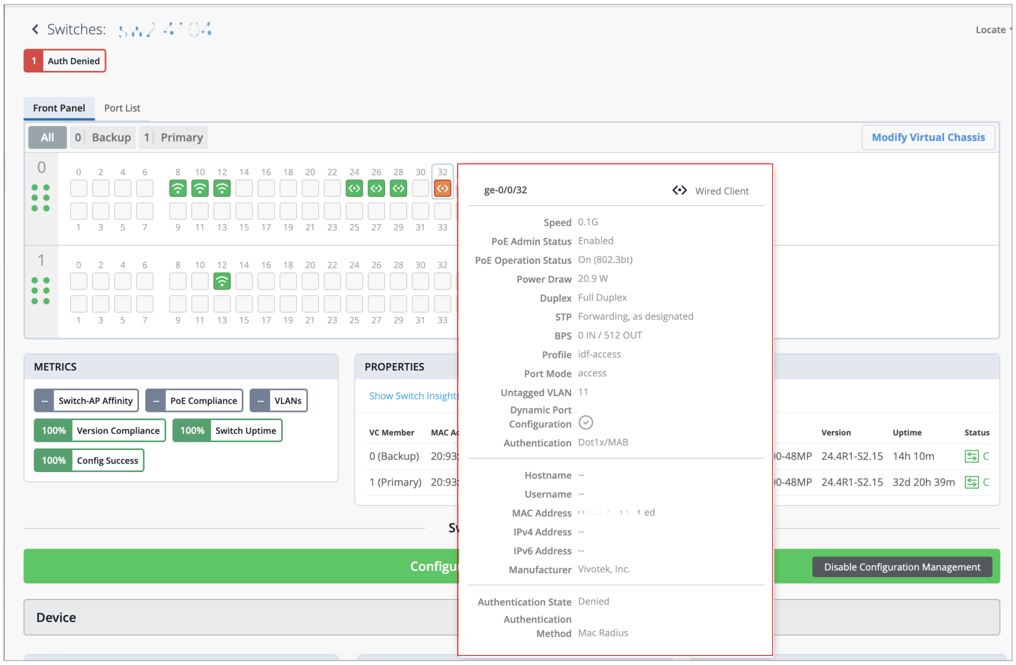

ge-0/0/32 — Wired Client

Speed: 0.1G
PoE Admin Status: Enabled
PoE Operation Status: On (802.3bt)
Power Draw: 20.9 W
Duplex: Full Duplex
STP: Forwarding, as designated
BPS: 0 IN / 512 OUT
Profile: idf-access
Port Mode: access
Untagged VLAN: 11
Dynamic Port Configuration: ✓
Authentication: Dot1x/MAB

Hostname: --
Username: --
MAC Address: [redacted] ed
IPv4 Address: --
IPv6 Address: --
Manufacturer: Vivotek, Inc.

Authentication State: Denied
Authentication Method: Mac Radius

VC Member	MAC A		Version	Uptime	Status
0 (Backup)	20:93	0-48MP	24.4R1-S2.15	14h 10m	C
1 (Primary)	20:93	0-48MP	24.4R1-S2.15	32d 20h 39m	C

To get a more detailed view of what's going on with a port, click that port to select it. When you select a port from the front panel, the following happens:

- If you select multiple ports, the configuration page displays the Port Configuration, Networks, and Port Profiles tiles with the settings applied to the selected ports. You can make configuration adjustments to the ports from these tiles. If the selected ports do not have any configuration, the Port Configuration, Networks, and Port Profiles tiles are displayed without any data.

- If you select only a single port, the configuration page additionally displays port Statistics and Wired Client insights. From the Wired Client insights tile, you can access the connected devices.

From the detailed view of the port, you can also bulk edit the port configuration, perform cable tests, bounce (restart) the port in case you encounter any issues with it, and clear MAC addresses stored through persistent MAC learning.

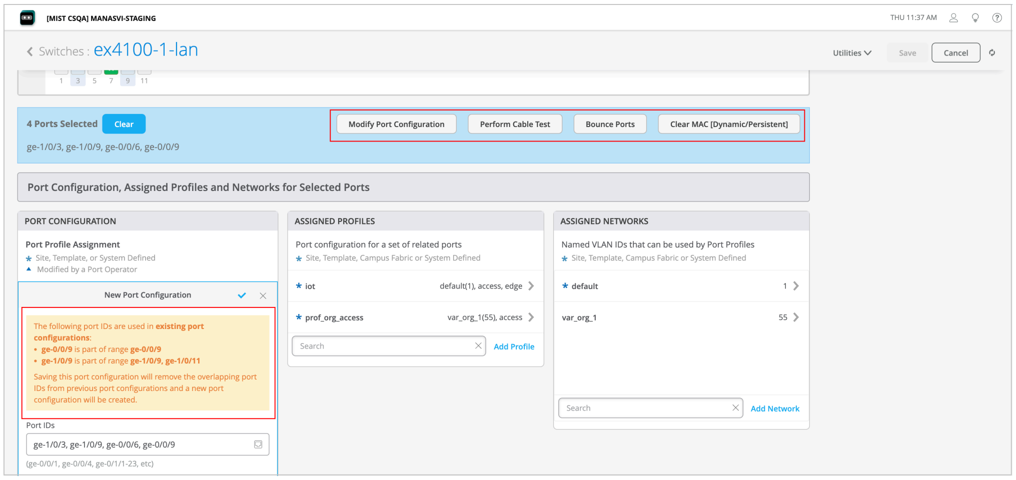

4 Ports Selected

ge-1/0/3, ge-1/0/9, ge-0/0/6, ge-0/0/9

Port Configuration, Assigned Profiles and Networks for Selected Ports

PORT CONFIGURATION

Port Profile Assignment
* Site, Template, or System Defined
▲ Modified by a Port Operator

The following port IDs are used in **existing port configurations**:

- **ge-0/0/9** is part of range **ge-0/0/9**
- **ge-1/0/9** is part of range **ge-1/0/9, ge-1/0/11**

Saving this port configuration will remove the overlapping port IDs from previous port configurations and a new port configuration will be created.

Port IDs: ge-1/0/3, ge-1/0/9, ge-0/0/6, ge-0/0/9
(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

ASSIGNED PROFILES

Port configuration for a set of related ports
* Site, Template, Campus Fabric or System Defined

* iot	default(1), access, edge
* prof_org_access	var_org_1(55), access

ASSIGNED NETWORKS

Named VLAN IDs that can be used by Port Profiles
* Site, Template, Campus Fabric or System Defined

* default	1
var_org_1	55

To bulk edit your port configuration or override any existing port configuration on a switch, select the ports to be configured from the Front Panel on the switch details page, and click **Modify Port Configuration** (shown in the above picture). In case the selected ports are already part of an existing port range configuration, a warning message indicating the same is displayed. When you save the new configuration, it will replace the existing configuration for the selected ports. Previously, you had to manually remove any port configuration overlap before you carry out a bulk port edit.

Sensors

Juniper switches include a variety of sensors to measure CPU, memory, temperature, power draw, power supply, and the cooling fans for each tray. Red/Yellow/Green indicators show the overall health, or you can hover over an indicator to show a detailed list. Red indicates a failure, or a value is in the critical range. Yellow is a warning, typically because a value is approaching a critical level. Greens means OK. Figure 1 shows the temperature for each monitored component on the switch.

Figure 23: Component Status

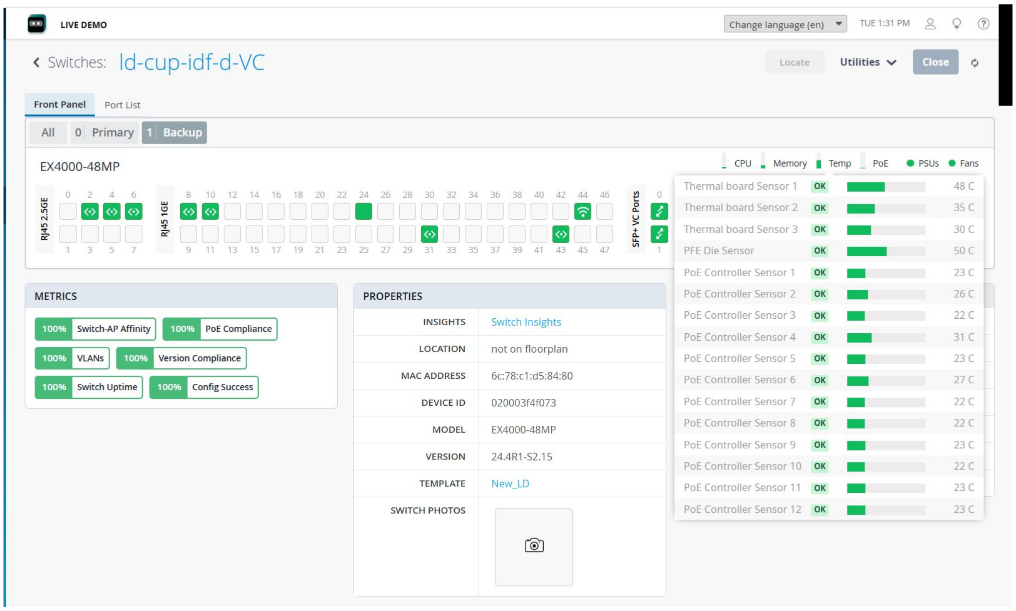

Thermal board Sensor 1	OK	48 C
Thermal board Sensor 2	OK	35 C
Thermal board Sensor 3	OK	30 C
PFE Die Sensor	OK	50 C
PoE Controller Sensor 1	OK	23 C
PoE Controller Sensor 2	OK	26 C
PoE Controller Sensor 3	OK	22 C
PoE Controller Sensor 4	OK	31 C
PoE Controller Sensor 5	OK	23 C
PoE Controller Sensor 6	OK	27 C
PoE Controller Sensor 7	OK	22 C
PoE Controller Sensor 8	OK	22 C
PoE Controller Sensor 9	OK	23 C
PoE Controller Sensor 10	OK	22 C
PoE Controller Sensor 11	OK	23 C
PoE Controller Sensor 12	OK	23 C

METRICS

100% Switch-AP Affinity
100% PoE Compliance
100% VLANs
100% Version Compliance
100% Switch Uptime
100% Config Success

PROPERTIES

INSIGHTS	Switch Insights
LOCATION	not on floorplan
MAC ADDRESS	6c:78:c1:d5:84:80
DEVICE ID	020003f4f073
MODEL	EX4000-48MP
VERSION	24.4R1-S2.15
TEMPLATE	New_LD
SWITCH PHOTOS	

Likewise, you can hover over the PSU indicator to see the status of each power supply on the switch. Green indicates the power supply is present and working properly. Red indicates insufficient power, or failure. Grey means the power supply is not present. Yellow indicates the power supply is initializing.

Live Port Traffic and Device Processes

Users with Super User or Network Admin privileges can get real-time statistics for a given switch interface by selecting the switch and then drilling down on the port. Traffic statistics include input and output, L2 and L3 errors, and BUM traffic, and are available on all EX series and QFX series switches with an active Wired Assurance license.

Figure 24: Live Traffic Counters

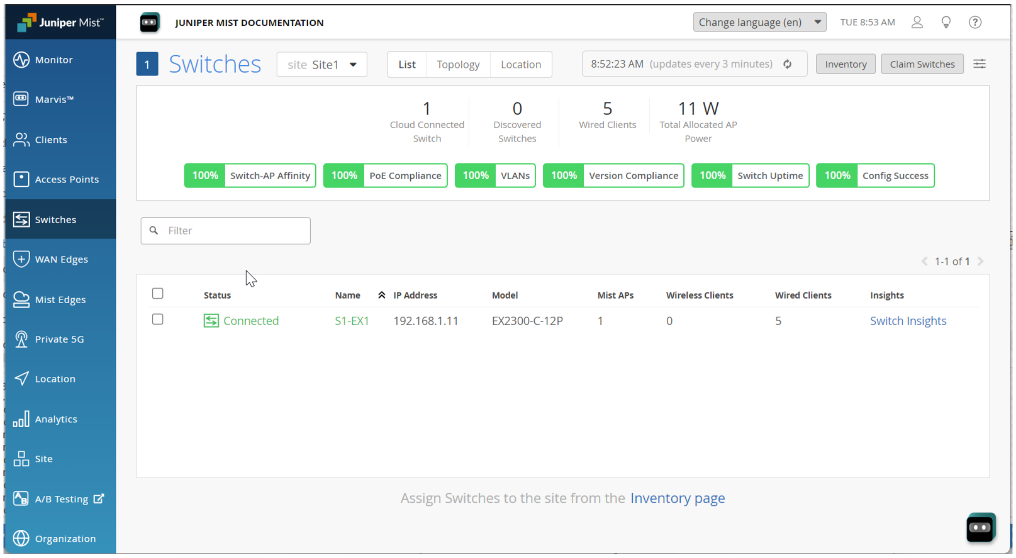

1 Cloud Connected Switch | 0 Discovered Switches | 5 Wired Clients | 11 W Total Allocated AP Power

100% Switch-AP Affinity | 100% PoE Compliance | 100% VLANs | 100% Version Compliance | 100% Switch Uptime | 100% Config Success

Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	Insights
Connected	S1-EX1	192.168.1.11	EX2300-C-12P	1	0	5	Switch Insights

Assign Switches to the site from the Inventory page

Viewing Switch Processes

You can get real-time statistics for the processes running on a given switch by selecting the switch and then drilling down to the switch Insights page, then clicking the **View Live Process Detail** button.

Figure 25: Live Process Detail

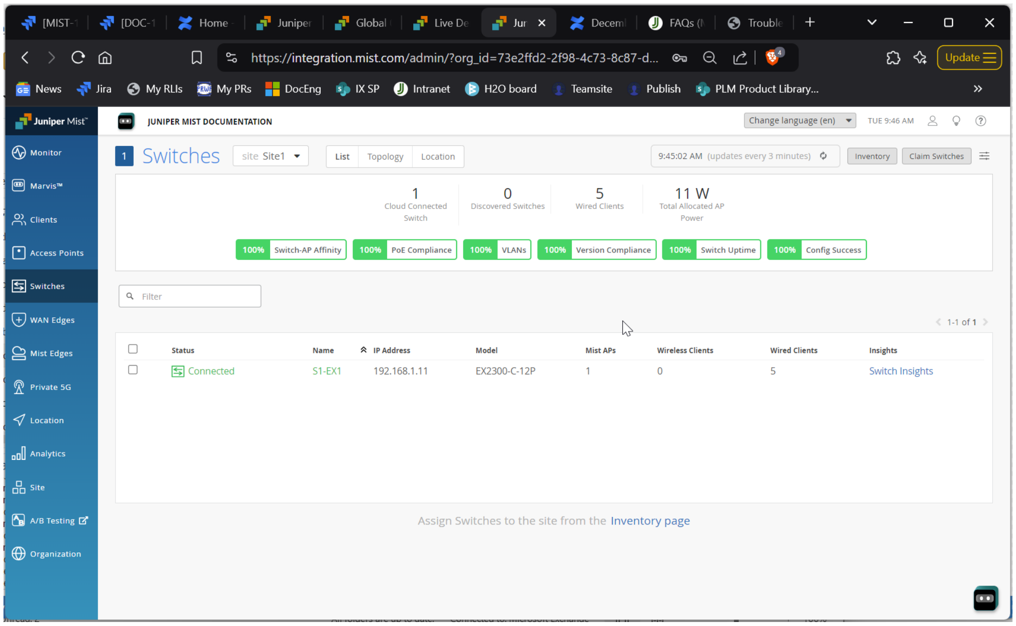

https://integration.mist.com/admin/?org_id=73e2ffd2-2f98-4c73-8c87-d...

1 Cloud Connected Switch | 0 Discovered Switches | 5 Wired Clients | 11 W Total Allocated AP Power

100% Switch-AP Affinity | 100% PoE Compliance | 100% VLANs | 100% Version Compliance | 100% Switch Uptime | 100% Config Success

Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	Insights
Connected	S1-EX1	192.168.1.11	EX2300-C-12P	1	0	5	Switch Insights

Assign Switches to the site from the Inventory page

Port List

If you want to see a list of all the ports, just click the **Port List** tab right beside the **Front Panel** tab.

In this Port List view, you'll find a wealth of information about each port. You can see the port name, its status, the clients connected to it, the power draw, port profile, port mode, port speed, and even the amount of data transmitted and received. You can also access each connected client by clicking the client name hyperlink.

By default, Port List displays a select set of columns. You can add more columns to the view from the list of available columns. To do that, go to Table Settings by clicking the hamburger menu in the upper right of the page. From the Table Settings page, select a check box to display the corresponding column (see the image below).

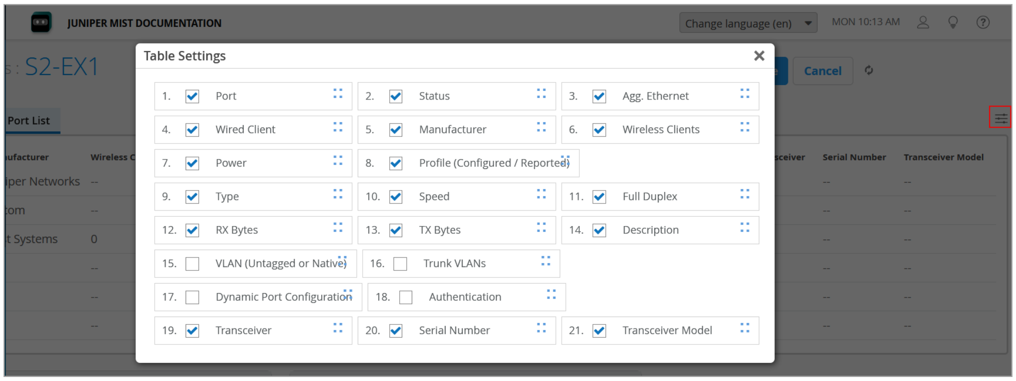

Table Settings

1.	☑	Port	2.	☑	Status	3.	☑	Agg. Ethernet
4.	☑	Wired Client	5.	☑	Manufacturer	6.	☑	Wireless Clients
7.	☑	Power	8.	☑	Profile (Configured / Reported)			
9.	☑	Type	10.	☑	Speed	11.	☑	Full Duplex
12.	☑	RX Bytes	13.	☑	TX Bytes	14.	☑	Description
15.	□	VLAN (Untagged or Native)	16.	□	Trunk VLANs			
17.	□	Dynamic Port Configuration	18.	□	Authentication			
19.	☑	Transceiver	20.	☑	Serial Number	21.	☑	Transceiver Model

Switch Insights

If you want to gain valuable insights into your switch, the events such as switch configuration changes, performance, and connected endpoints, visit the Switch Insights page by clicking the **Switch Insights** hyperlink on the **Properties** tile.

> **NOTE**: You can also go to the Switch Insights page by using the hyperlink on the main Switches page.

Switch Insights gives you a bird's-eye view of all the switch events that have taken place. It's like a log of configuration changes, software updates, and system alarms.

Switch Insights also lets you track some important metrics like CPU and memory utilization of the switch. You can also dive into details about BGP neighbors (applicable to campus fabrics). You can also view traffic patterns, port errors, and power draw at the switch or port level.

The key sections on the Switch Insights page are the following:

- Switch Events—Provides a log of all the switch events reported. You can filter good, neutral, and bad events.
- Table Capacity—Provides the following indicators:
  - MAC Address Table—Displays the percentage of the MAC address table capacity used. The MAC address table contains MAC Address-interface bindings associated with each VLAN.
  - ARP Table—Displays the percentage of Address Resolution Protocol (ARP) table capacity used. The ARP table contains the learned MAC Address-IP bindings of the devices connected to the network.
  - Route Summary—Displays the percentage of routing table capacity used.
  - EVPN Database—Available for switches that are part of an EVPN topology.

  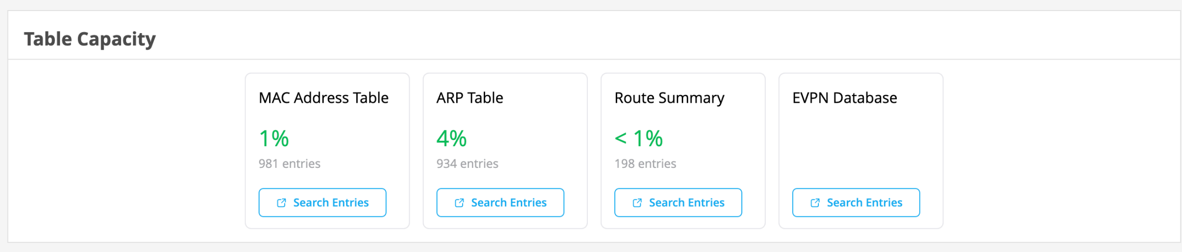

  Table Capacity

MAC Address Table	ARP Table	Route Summary	EVPN Database
1%	4%	< 1%	
981 entries	934 entries	198 entries	

  You can click the **Search Entries** button under each metric to open a shell view in a new window where you can search for entries after specifying filters. You also have the option to refresh and clear the entries displayed. Clicking **Refresh** on the upper right of the window provides a continuous display of the entry every three seconds for a total of 30 seconds. To stop the refresh before the 30-second timer is complete, close the window or click another table. Clicking the **Clear *Entry*** button, which is available only for MAC and ARP table, clears the respective entry from the table. You also have the option to clear the buffer on the screen by clicking **Clear Screen** at the lower left of the window.

- Switch Charts—A set of charts providing insights into key metrics.

  Use the drop-down menu at the top of this section to select all ports or a specific port.

  Charts include:

  - CPU Utilization—Percent of CPU in use across the selected time period. Hover your mouse over any point in time to see details below the chart.

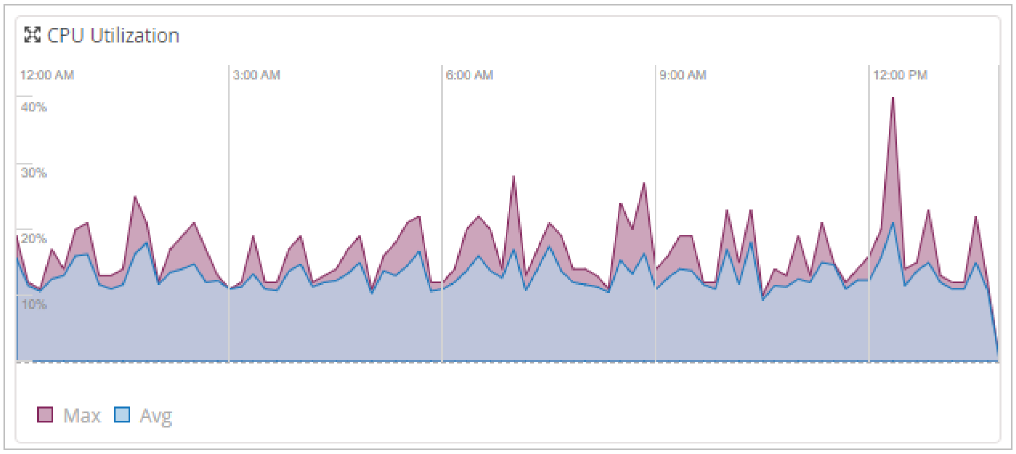

- Memory Utilization—Percent of memory in use across the selected time period. Hover your mouse over any point in time to see details below the chart.

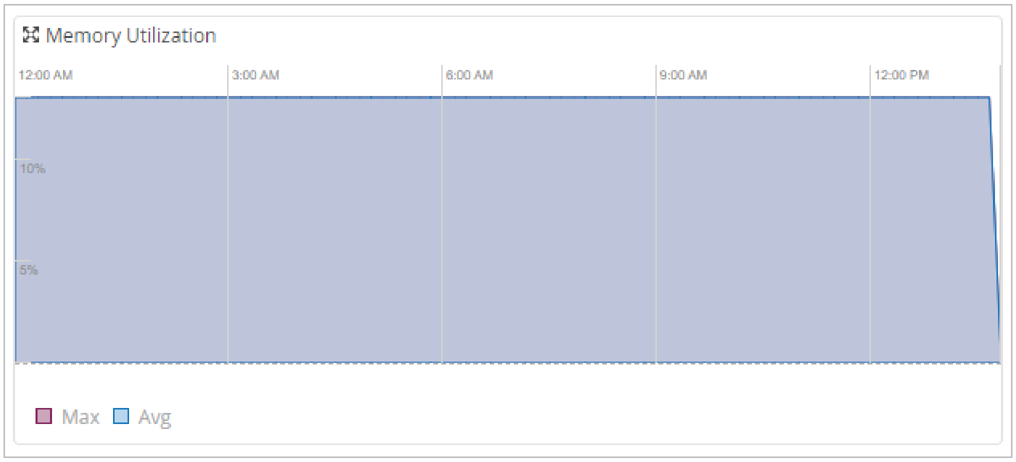

- Bytes—Number of bytes transmitted (TX) and received (RX) across the selected time period. Hover your mouse over any point in time to see details in the white bar at the center of the chart.

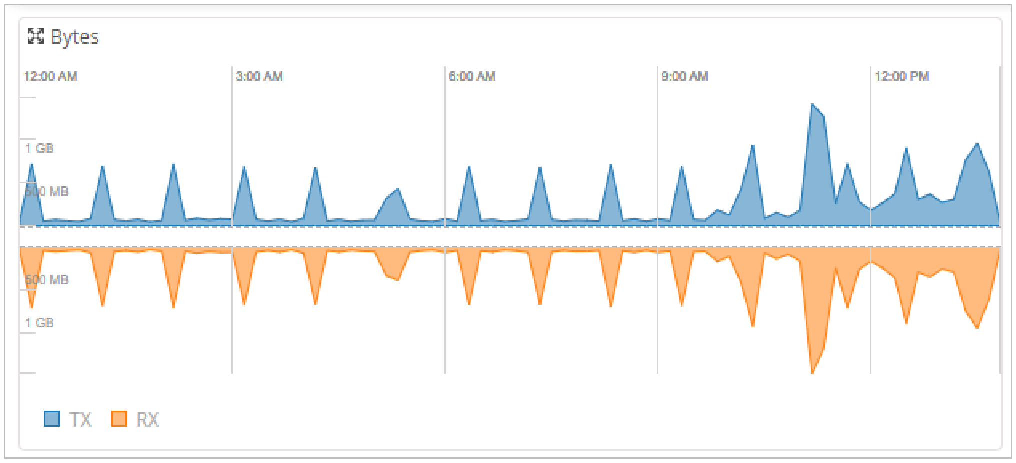

- Data Rate—Transmit (TX) and receive (RX) rate (in mbps) across the selected time period. Hover your mouse over any point in time to see details in the white bar at the center of the chart.

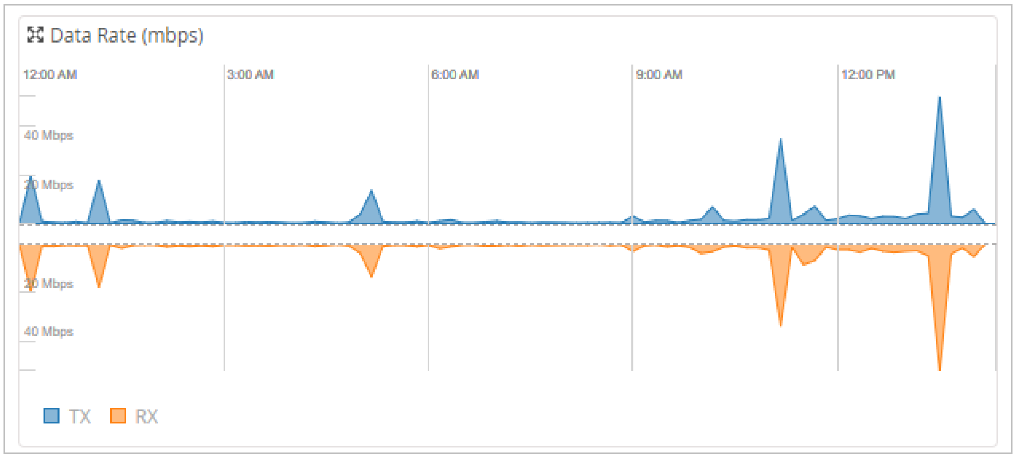

- TX/RX Packets—Number of packets transmitted (TX) and received (RX) across the selected time period. Color coding shows the different types of packets (multicast, broadcast, and unicast). Hover your mouse over any point in time to see details in a pop-up description.

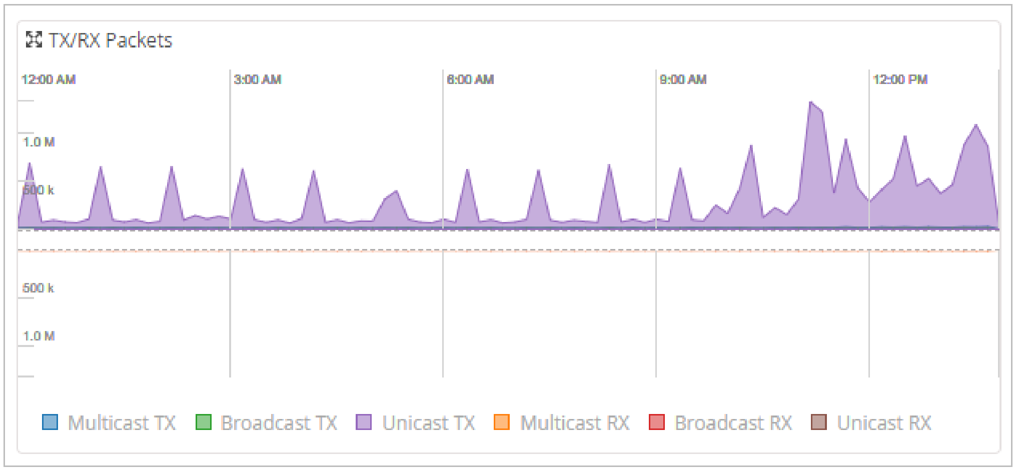

- Port Errors—Number of transmit (TX) and receive (RX) errors across the selected time period. Hover your mouse over any point in time to see details below the chart.

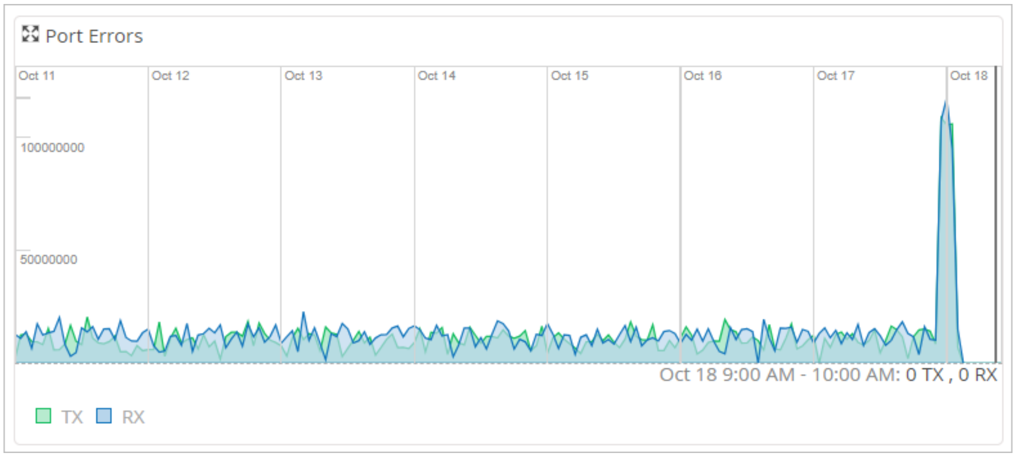

- Power Draw—Maximum and minimim Watts drawn across the selected time period. Hover your mouse over any point in time to see details below the chart.

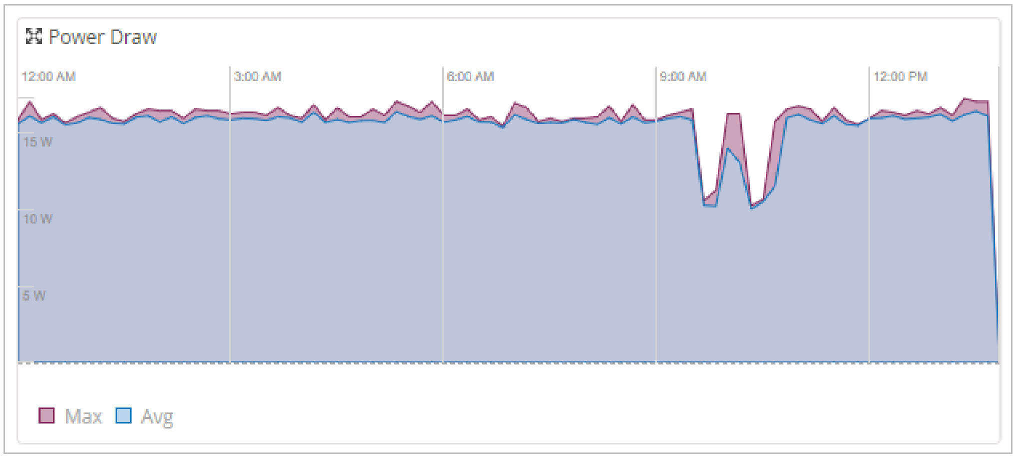

- Switch Ports—A detailed port list. You can see all the ports along with information about the connected endpoints.
- Current Switch Properties—Displays the switch properties along with its connection status. The properties include the switch location, MAC address, switch model, Junos version, BIOS version, U-boot version, and POE version.

Video: Switch Insights Overview

Metrics

The **Metrics** section on the switch details page helps you track the performance of your switches against specific compliance parameters, and identify if there are any areas that need attention. You'll find a bunch of important compliance parameters that are being monitored. For example, you can keep an eye on switch-AP affinity, version compliance, PoE compliance, and switch uptime. When you click each metric, you'll be taken to a detailed view that provides more information.

For more information, see "Switch Metrics" on page 293.

Properties

The Properties tile displays the switch properties that include the device location, MAC address, device ID, device model, and the Junos version on the device. You can access Switch Insights from the Properties tile.

Statistics

This tile displays the switch connection status along with other details such as IP address, number of connected APs, uptime, and configuration status. The Last Config field on the Statistics tile shows the configuration change timestamp immediately after a user makes a configuration change to the switch.

Switch Utilities

The switch details page provides troubleshooting and testing tools to help you get to the root of any issue. To access these tools, click the **Utilities** drop-down list in the upper right corner of the page.

For more information on switch utilities, see "Switch Utilities" on page 307.

Switch Configuration

If you need to take a closer look at the configuration or make changes to it, scroll down to the **Switch Configuration** section on the switch details page. This section shows all the configurations applied to the switch through the template linked to the site to which the switch is onboarded. It also shows additional switch-specific settings.

To learn more about the switch configuration templates and different configuration options, see "Switch Configuration" on page 22.

Switch Utilities

SUMMARY

Use the Switch Details page to check connectivity, access the remote shell, restart the switch, and take a recovery snapshot, and use other utilities,

The Switch Utilities on the switch details page help you troubleshoot and test your switch.

To access the switch utilities:

1. On the Mist portal, click **Switches** on the left pane.

2. Click a switch from the **List** tab to open the switch details page.

3. Select the utility or tool from the **Utilities** drop-down list on the upper right of the page.

The switch utilities include the following tools:

- **Testing tools**—Use the switch testing tools to check the switch connectivity and monitor the switch health. The following tools are available:

  - **Ping**—Check the host reachability and network connectivity. To run the test, specify a hostname and then click **Ping**. The ping utility does a look-up for the hostname in the default routing instance (inet.0). You can specify the number of packets to be transmitted in the test.

  - **Traceroute**—View the route that packets take to a specified network host. Use this option as a debugging tool to locate the points of failure in a network. To run the test, specify a hostname, port, and a timeout value, and then click **Traceroute**.

  - **Cable Test** —Run a **Cable Test** on a port to monitor the connection health of the cable on the specified port. To run the test, specify the interface name (example: ge-0/0/4) and then click **Cable Test**. This action runs a time domain reflectometry (TDR) diagnostic test on the specified interface and displays results. A TDR test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

  - **Bounce Port**—Restart any unresponsive ports on your switch. To restart a port, specify the interface name (example: ge-0/0/4) and then click **Bounce Port**.

- **Remote Shell**—Access the command line directly through the Mist portal. You can enter commands on the switch's CLI without making a physical connection to the console port or using SSH. You can download a log of your remote shell session by clicking the download button in the upper right portion of the Remote Shell session.

- **Send Switch Log to Mist**—When you experience an issue with a switch, use this option to securely send the switch logs to Mist. The Juniper Mist support team uses these logs to understand the issue and provide troubleshooting support. This action sends the following log files:
  - RSI
  - /var/log/messages
  - PHC (phone home client) logs
  - Log files for troubleshooting cloud connectivity
- **Reboot Switch**—Reboot the switch directly from the switch details page.
- **Upgrade Firmware**—Upgrade switches directly from the switch details page. Before performing an upgrade, ensure that the switch has enough storage and a stable SSH connectivity to Mist cloud. If the switch doesn't have enough space, the upgrade will fail. To free up the space, run the following command as a root user:

  ```
  user@switch01> start shell user root
  root@Switch01:RE:0% pkg setop rm previous
  root@Switch01:RE:0% pkg delete old
  ```

  For more information, see "Upgrade Junos OS on Switches" on page 153.
- **Create Template**—Create a new switch configuration template. See also: "Create a Switch Configuration Template" on page 22.
- **Snapshot Device**—Store a recovery snapshot in the OAM (Operations, Administrations and Maintenance) volume. The recovery snapshot in the OAM volume can be used to boot the switch in case of corrupt configuration or firmware update failure.
- **Download Junos Config**—Use this tool to download the switch's Junos configuration in a text file.
- **Replace Switch**—Replace a switch without disrupting the network services on your network topology. By default, this action copies the existing switch configuration to the new switch. You can also choose to discard specific configurations that you don't want to copy to the new switch. For more information, see "Replace a Switch" on page 176.

  > **NOTE**: If the new switch has a different number of ports than the switch being replaced, the port configuration is discarded. If the current switch template doesn't

cover the configuration requirement of your new switch, we recommend that you assign your site with a different template that covers the new switch. See "Configure Switches Using Templates" on page 21.

- **Sync Configuration**—Sometimes, the configuration push to a switch might fail for various reasons. In such cases, you can use this option to resend the configuration to the switch. The **Sync Configuration** operation will overwrite any configuration defined on the switch manually through the CLIs.

RELATED DOCUMENTATION

Switch Details | **294**

Wired Clients

SUMMARY

The wired clients page provides a list of devices connected to the switches in a site. You can choose which client parameters to display in the list and/or download a local copy.

You can a see a list of wired clients, including the current connection status, for a given site in the Wired Client page. By default, the page shows the MAC address, VLAN IDs, and provides a shortcut to the device Insights page.

Figure 26: The Wired Clients Page

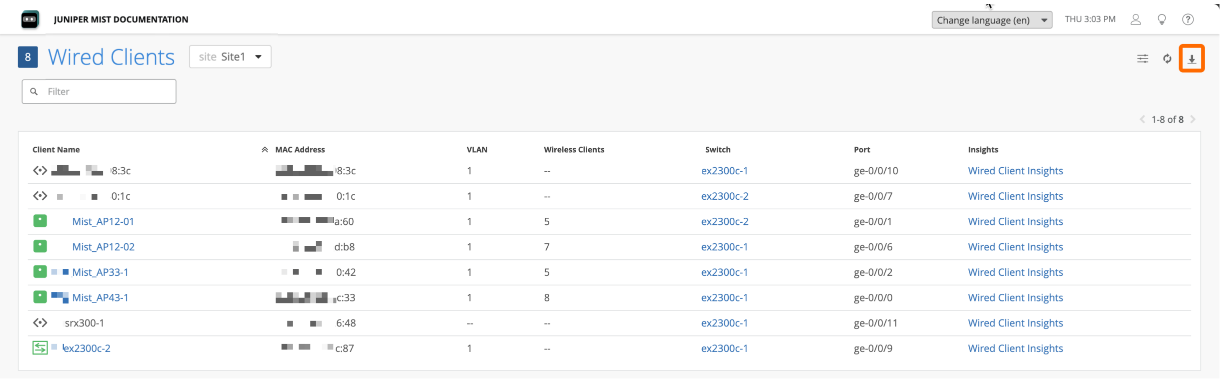

Client Name	MAC Address	VLAN	Wireless Clients	Switch	Port	Insights
[redacted]8:3c	[redacted]8:3c	1	--	ex2300c-1	ge-0/0/10	Wired Client Insights
[redacted]0:1c	[redacted]0:1c	1	--	ex2300c-2	ge-0/0/7	Wired Client Insights
Mist_AP12-01	[redacted]a:60	1	5	ex2300c-2	ge-0/0/1	Wired Client Insights
Mist_AP12-02	[redacted]d:b8	1	7	ex2300c-1	ge-0/0/6	Wired Client Insights
[redacted]Mist_AP33-1	[redacted]0:42	1	5	ex2300c-1	ge-0/0/2	Wired Client Insights
[redacted]Mist_AP43-1	[redacted]c:33	1	8	ex2300c-1	ge-0/0/0	Wired Client Insights
srx300-1	[redacted]6:48	--	--	ex2300c-1	ge-0/0/11	Wired Client Insights
[redacted]ex2300c-2	[redacted]c:87	1	--	ex2300c-1	ge-0/0/9	Wired Client Insights

Most of the fields available are not visible in the Wired Clients list view, because they make the table too wide to see without scrolling to the right. To display additional field in the list, click the Table Settings (☰) icon at the top of the page to select the columns you want to display.

Figure 27: Available Wired Clients Values

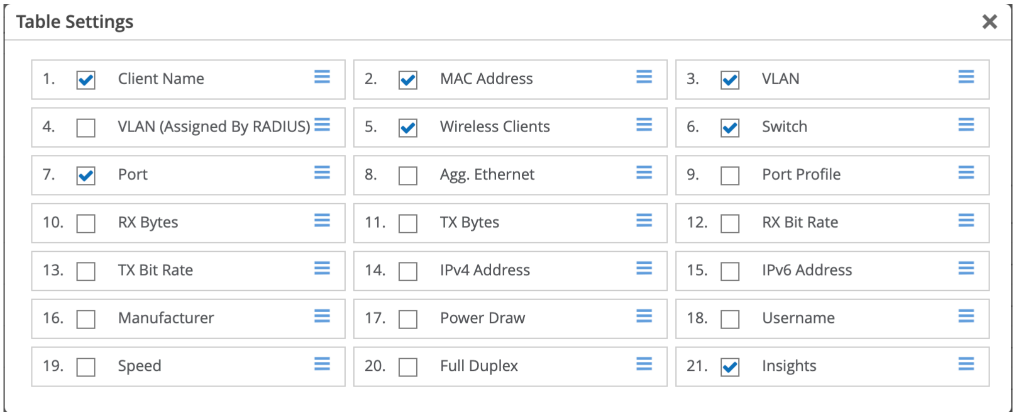

Table Settings

1. ☑ Client Name	2. ☑ MAC Address	3. ☑ VLAN
4. ☐ VLAN (Assigned By RADIUS)	5. ☑ Wireless Clients	6. ☑ Switch
7. ☑ Port	8. ☐ Agg. Ethernet	9. ☐ Port Profile
10. ☐ RX Bytes	11. ☐ TX Bytes	12. ☐ RX Bit Rate
13. ☐ TX Bit Rate	14. ☐ IPv4 Address	15. ☐ IPv6 Address
16. ☐ Manufacturer	17. ☐ Power Draw	18. ☐ Username
19. ☐ Speed	20. ☐ Full Duplex	21. ☑ Insights

For example, the following DHCP client fields are available for wired clients, but are not shown unless selected in the tables settings page:

- **Dynamic Filter**—This field indicates whether the client is included in any dynamic filtering rules or policies (these policies are configured on the switch but are typically assigned to clients dynamically through NAC or RADIUS).
- **RADIUS Returned VoIP VLAN**—Shows the VoIP VLAN information returned by the RADIUS server.
- **Auth Domain**—Indicates the authentication domain classification (e.g., VOIP, DATA).
- **DHCP FQDN** —(Requires DHCP Snooping Junos version 23.2 or later on the client.)

  Displays the fully qualified domain name.

- **DHCP Hostname**—(Requires DHCP Snooping Junos version 23.2 or later on the client.)
- **DHCP Vendor Class Identifier**—(Requires DHCP Snooping Junos version 23.2 or later on the client.)
- **DHCP Client Identifier**—(Requires DHCP Snooping Junos version 23.2 or later on the client.)

The table describes the most common fields displayed for wired clients.

Table 23: The Wired Clients Page

Column Name	Details
Client Name	The client name is either a MAC address, the assigned name of a switch, AP, WAN edge, or Mist Edge device. The icon shown along with the name is the same one shown on the front panel display of the "Switch Details" on page 294 page.
MAC Address	Displays the MAC address for the client connected to that switch port.
VLAN	Shows the assigned VLAN for the switch port.
Wireless Clients	When an AP is attached to the switch, this column shows the count of wireless clients attached to the AP.
Switch	When there are multiple switches in a site, this column lists the name of the switch to which the wired device is attached.
Port	This lists the port on the switch to which the client devices is attached.
Insights	This is a link to the Wired Clients Insights page for the specific wired client on that row of the list.

6

CHAPTER

Troubleshooting

SUMMARY

Use the information in this chapter to troubleshoot issues on your wired network.

IN THIS CHAPTER

What Do You Want to Do?

Table 24: Top Tasks

If you want to...	Use these resources:
Assess network health. *Use the service-level expectation (SLE) dashboards to assess the user experience and resolve any issues proactively.*	"Wired SLEs" on page 313
Troubleshoot with Marvis. *If you have a Marvis subscription, troubleshoot with the Actions dashboard and the conversational assistant.*	"Troubleshoot with Marvis" on page 340
Investigate specific issues. *Use the other topics in this chapter to get help with other issues.*	Troubleshooting topics in this chapter
Explore other aspects of troubleshooting and AI-driven operations. *Take a deeper dive into the various dashboards, tools, and Marvis features that support AI-driven operations.*	Juniper Mist AI-Drive Operations Guide

Wired SLEs

SUMMARY

Use the wired service-level experience (SLE) dashboard to assess the service levels for user-

IN THIS SECTION

impacting factors such as throughput, connectivity, and switch health.

Overview

IN THIS SECTION

Juniper Mist™ cloud continuously collects network telemetry data and uses machine learning (ML) to analyze the end-user experience. This service efficiently collects and analyzes data your entire network, whether you have hundreds or thousands of ports.

You can access this information through the Juniper Mist wired service-level expectation (SLE) dashboards, which help you assess the network's user experience and resolve any issues proactively. It's not merely a matter of devices or links being up or down—it's the quality of the client experience.

For the wired network, the two burning questions are:

- Are clients able to connect?
- Are clients able to pass traffic after connecting?

The wired SLE dashboards show the user experience of the wired clients on your network at any given point in time. You can use these interactive dashboards to measure and manage your network proactively by identifying any user pain points before they become too big of an issue.

Finding the Wired SLEs Dashboard

To find the Wired SLEs dashboard, select **Monitor** > **Service Levels** from the left menu, and then click the **Wired** button.

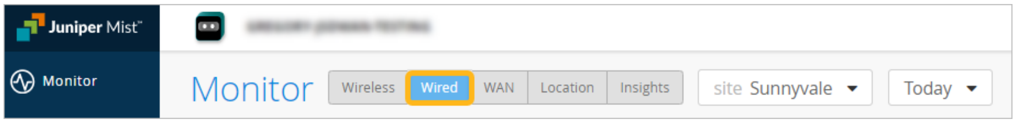

NOTE: The buttons appear only if you have the required subscriptions. For information about these requirements, see the Juniper Mist AI-Native Operations Guide.

Root Cause Analysis for the Wired Successful Connect SLE

After you click a classifier in the SLE block, you'll see the Root Cause Analysis page. Click classifiers and sub-classifiers to view timeline and scope information in the lower half of the screen.

NOTE: The information in the lower half of the screen depends on what you've selected at the top.

Useful tabs in the lower half of the screen are:

- Timeline—See exactly when the issues occurred.
- Distribution—See which VLANs were affected.
- Affected Items—See which interfaces and clients were affected and how much each one contributed to the overall impact. Also see the individual failure rate for each interface or client.

Let's look at an example for the Successful Connect SLE. By clicking options at the top of the page, you can drill down from the SLE to classifiers and sub-classifiers. The lower half of the page shows information relevant to these selections.

By selecting the **Affected Items** tab and then clicking the **Interfaces** option on the left, we see the interfaces that were unable to connect due to incorrect credentials.

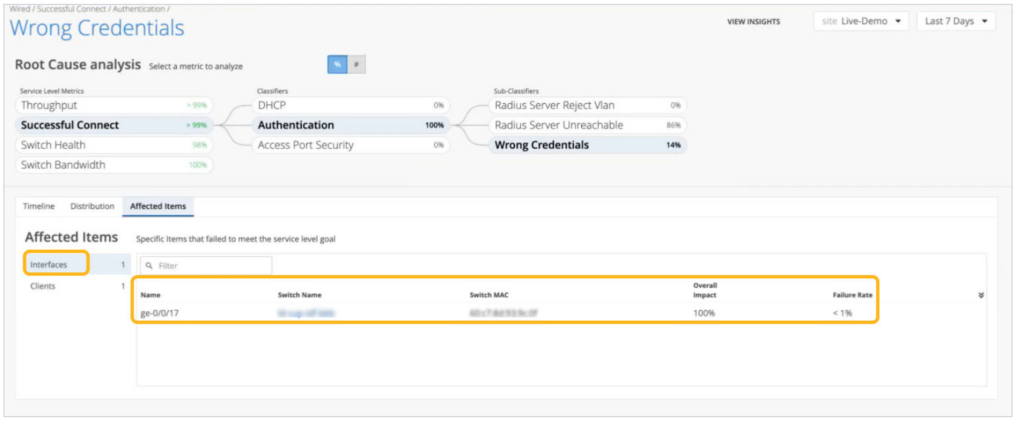

Wrong Credentials

Root Cause analysis Select a metric to analyze

Service Level Metrics		Classifiers		Sub-Classifiers	
Throughput	> 99%	DHCP	0%	Radius Server Reject Vlan	0%
Successful Connect	> 99%	Authentication	100%	Radius Server Unreachable	86%
Switch Health	98%	Access Port Security	0%	Wrong Credentials	14%
Switch Bandwidth	100%				

Affected Items Specific Items that failed to meet the service level goal

Interfaces	1
Clients	1

Name	Switch Name	Switch MAC	Overall Impact	Failure Rate
ge-0/0/17	[redacted]	[redacted]	100%	< 1%

By clicking the **Clients** tab on the left, we now see the affected clients.

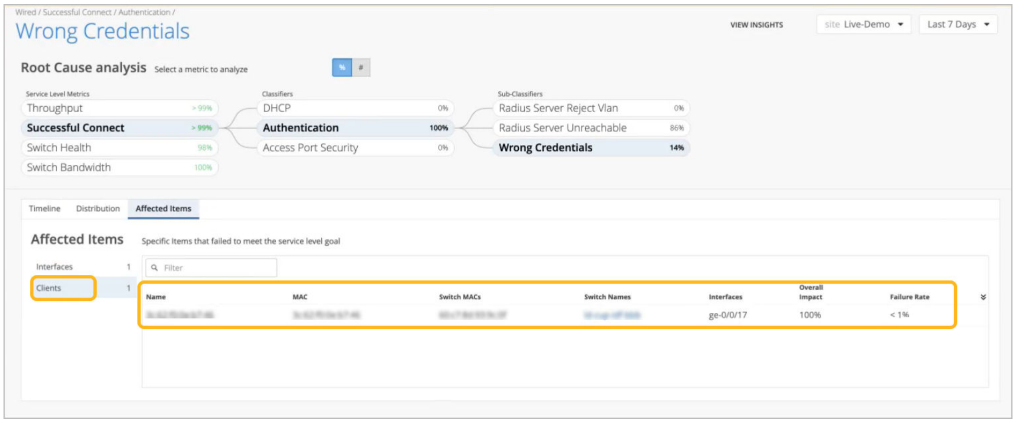

Wired / Successful Connect / Authentication /

Wrong Credentials

VIEW INSIGHTS | site Live-Demo ▾ | Last 7 Days ▾

Root Cause analysis Select a metric to analyze % | #

Service Level Metrics		Classifiers		Sub-Classifiers	
Throughput	> 99%	DHCP	0%	Radius Server Reject Vlan	0%
Successful Connect	> 99%	**Authentication**	100%	Radius Server Unreachable	86%
Switch Health	98%	Access Port Security	0%	**Wrong Credentials**	14%
Switch Bandwidth	100%				

Timeline | Distribution | **Affected Items**

Affected Items Specific Items that failed to meet the service level goal

Interfaces 1
Clients 1

Filter

Name	MAC	Switch MACs	Switch Names	Interfaces	Overall Impact	Failure Rate
[redacted]	[redacted]	[redacted]	[redacted]	ge-0/0/17	100%	< 1%

TIP:

- **Overall Impact** is the percentage that a client or interface contributed to *all* issues for the selected sub-classifier. For example, it can show if a client account for 20 percent or 90 percent of the issues.

- **Failure Rate** is the impact of this issue on this interface or client. For example, it can show if an interface was unsuccessful on 20 percent or 90 percent of connection attempts.

- To see more details, click the hyperlinks in the table to go to the Insights page, where you can see all client and switch events.

Wired Assurance: Day 2 - Wired SLEs Video Overview

Video: Wired Assurance: Day 2 - Wired Service Level Expectations (SLEs)

Wired SLE Blocks

As shown in the following example, each SLE block provides valuable information.

- At the left, you see that this SLE has an 89 percent success rate.

  At the center, the timeline shows variations across the time period. You can hover your mouse pointer over any point to see the exact time and SLE outcome.

  At the right, the classifiers show what percentage of issues were attributable to each root cause. In this example, 100 percent of the issues were attributed to Network.

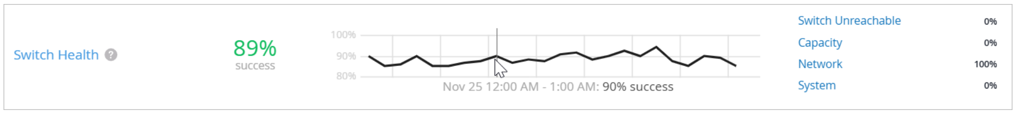

If you click a classifier, you'll see more information on the Root Cause Analysis page. Most classifiers have sub-classifiers for greater insight into the exact causes of issues.

The following table provides more information about the wired SLEs and classifiers.

Table 25: Wired SLE Descriptions

SLE	SLE Description	Classifier	Classifier Description
Successful Connect	Juniper Mist monitors client connection attempts and identifies failures. The source of data is 802.1X events on the switch. This SLE helps you to assess the impact of these failures and to identify the root causes to address. This SLE is available if you use 802.1X on the wired network to authenticate clients or if you have DHCP snooping configured. You cannot set the threshold for this SLE. It's assumed that you want 100 percent successful connects and consider any unsuccessful connect as a critical issue to track.	DHCP	Client connections that fail to reach the bound state within a minute. This classifier is available only when DHCP snooping is enabled in the port profile. DHCP snooping might not always work well with endpoints that have static IPs.
		Authentication	Events when a client failed to authenticate. Sub-Classifiers: • RADIUS Server Reject VLAN—Couldn't authenticate to the specified VLAN. • Wrong Credentials—The credentials weren't valid. • RADIUS Server Unreachable—The RADIUS server was down.

Table 25: Wired SLE Descriptions *(Continued)*

SLE	SLE Description	Classifier	Classifier Description
		Access Port Security	Client connection failures caused by access port security issues. Based on the security features configured in your port profiles, this classifier is triggered as security events occur. Sub-Classifiers: • BPDU-Guard—Detects connection failures because of the BPDU guard configuration on the switch port. This feature is important to prevent looping, as when a switch is connected to a switch. To enable this feature, go to the port profile, and enable STP Edge. • MAC Limit—Detects connection failures reported when a client exceeds the MAC limit configured on the switch port. For example, you might configure your port profile with a MAC limit of 2 if you have an outdoor security camera or public address system and want to prevent other devices from connecting to that

Table 25: Wired SLE Descriptions *(Continued)*

SLE	SLE Description	Classifier	Classifier Description
			port. If someone unplugs your camera and attempts to connect their own device, the MAC limit would be reached, and this event would be reflected by the MAC limit classifier. • Dynamic ARP Inspection—Identifies client connection failures when a port drops invalid Dynamic ARP Inspection packets. This security feature prevents people from snooping for someone else's ARP address to gain access. Requires enabling ARP Inspection in the DHCP Snooping section of the port profile. • Rogue DHCP Server—Identifies client connection failures caused by a rogue DHCP server event. This could be an event where an untrusted port drops traffic from DHCP servers to block unauthorized servers. Enabling this feature can prevent rogue devices from

Table 25: Wired SLE Descriptions *(Continued)*

SLE	SLE Description	Classifier	Classifier Description
			connecting. This classifier shows any such attempts that occur. Requires enabling DHCP snooping in the port profile.
Throughput	This SLE represents the ability of wired users to pass traffic without impedance. You cannot set the threshold for this SLE. It's assumed that you want 100 percent of traffic to pass without impedance and consider any impedance as a critical issue to track.	Storm Control	Events when storm control level was exceeded and packets were dropped. Available only if you've enabled Storm Control in the port profile (recommended).

Table 25: Wired SLE Descriptions ***(Continued)***

SLE	SLE Description	Classifier	Classifier Description
		Interface Anomalies	Events when devices were powered up but could not pass traffic. Sub-Classifiers: • **Cable Issues**—This sub-classifier shows the user minutes affected by faulty cables in the network. Cable issues can cause a high failure rate on an interface or client device. • **Negotiation Failed**—This sub-classifier identifies bad user minutes caused by issues such as incomplete negotiation, duplex conflict, and latency. • **MTU Mismatch**—This sub-classifier identifies issues where MTU size is mismatched somewhere along the packet's path (any MTU mismatch along the path will result in discarded or fragmented packets). The information for this SLE comes from the switch; each input error or MTU error contributes to a bad

Table 25: Wired SLE Descriptions *(Continued)*

SLE	SLE Description	Classifier	Classifier Description
			user minute under this sub-classifier.
Switch Bandwidth	Juniper Mist™ measures the available bandwidth on your network based on the queued packets and dropped packets for each configured queue. A pattern of low success rates can indicate a need for more wired bandwidth. You can click the **Settings** button to set the percentage to use as the success threshold for this SLE. This percentage represents the total_DropppedPackets as a portion of total_QueuedPackets.	Congestion	Heavy congestion causing dropped packets (TxDrops) when the input queue (buffer) fills up. Triggered by considering these ratios: • TxDrops to TxPackets—Total transmitted bytes dropped to total packets transmitted. • Txbps to Link speed—Total bytes transmitted per second to link speed. • RxSpeed to Link speed—Total bytes received per second to link speed.

Table 25: Wired SLE Descriptions *(Continued)*

SLE	SLE Description	Classifier	Classifier Description
		Congestion Uplink	High congestion on uplinks with these uplink port characteristics: • Has a switch or a router as an LLDP neighbor • Is a Spanning Tree Protocol (STP) root port • Has a higher number of transmitted and received packets compared to the other ports • Experiencing congestion due to aggregated Ethernet links and module ports
		Bandwidth Headroom	High bandwidth usage.
Switch Health	Juniper Mist™ monitors your switches' operating temperature, power consumption, CPU, and memory usage. Monitoring switch health is crucial because issues such as high CPU usage can directly impact connected clients. For example, if CPU utilization spikes to 100 percent, the connected APs might lose	Switch Unreachable	Poor switch-to-cloud connectivity. The switch might be down, or the connection might be severed.

Table 25: Wired SLE Descriptions *(Continued)*

SLE	SLE Description	Classifier	Classifier Description
	connectivity, affecting the clients' experience.	Capacity	Usage exceeding 80 percent. High usage can indicate that the switch is dealing with more requests that it can optimally handle. Sub-Classifiers indicate usages exceeding 90 percent of the relevant table capacity: • ARP Table • Route Table • MAC Table
		Network	Lower than expected throughput due to uplink capacity limitations. Based on the round-trip time (RTT) value of packets sent from the switch to the Mist cloud. Sub-Classifiers: • **WAN Latency**—Based on the average value of RTT over a period of time. • **WAN Jitter**—Calculated by comparing the standard deviation of RTT within a small period with the overall deviation of RTT over a longer period.

Table 25: Wired SLE Descriptions *(Continued)*

SLE	SLE Description	Classifier	Classifier Description
		System	Issues on the switch that can impact user experiences Sub-Classifiers: • **CPU**—Utilization above 90 percent • **Memory**—Utilization above 80 percent • **Temp**—Temperature above or below the specified operating range • **Power**—Consumption above 90 percent of the available power

Troubleshoot Your Switch Connectivity

SUMMARY

Follow these guidelines to check traffic statistics and troubleshoot problems with switch connectivity.

You can get real-time statistics for a given switch interface by selecting the switch and clicking the **Live Traffic Counters** button, or drilling down on a specific interface or virtual chassis member and then doing the same. Traffic statistics include input and output, L2 and L3 errors, and BUM traffic.

Figure 28: Live Traffic Counters

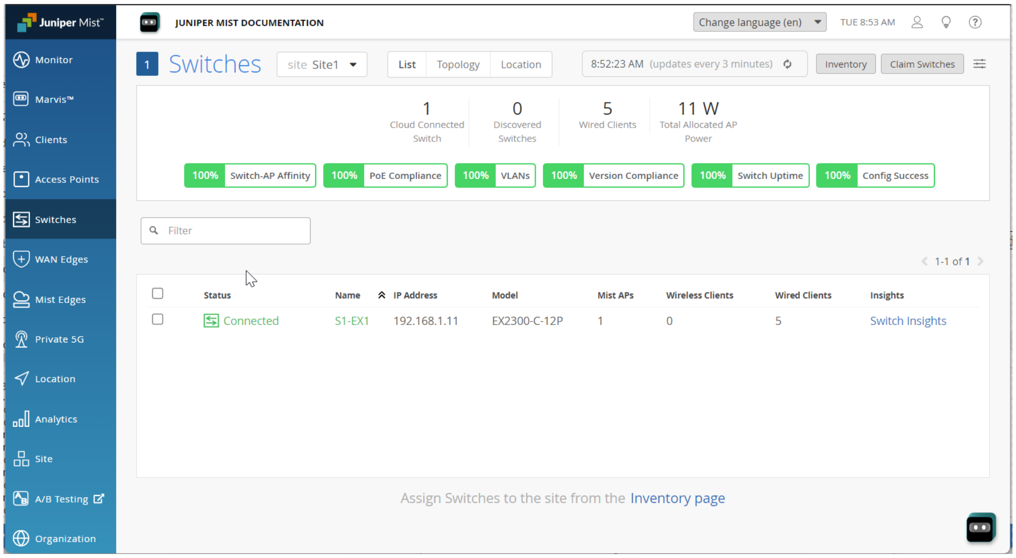

1 Cloud Connected Switch | 0 Discovered Switches | 5 Wired Clients | 11 W Total Allocated AP Power

100% Switch-AP Affinity | 100% PoE Compliance | 100% VLANs | 100% Version Compliance | 100% Switch Uptime | 100% Config Success

	Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	Insights
□	Connected	S1-EX1	192.168.1.11	EX2300-C-12P	1	0	5	Switch Insights

Assign Switches to the site from the Inventory page

View Switch Processes

Real-time statistics for the processes running on a given switch are available by selecting the switch and then drilling down to the switch Insights page. Click the **View Live Process Detail** button. Supports virtual chassis and virtual machines in addition to physical devices.

Figure 29: Live Process Detail

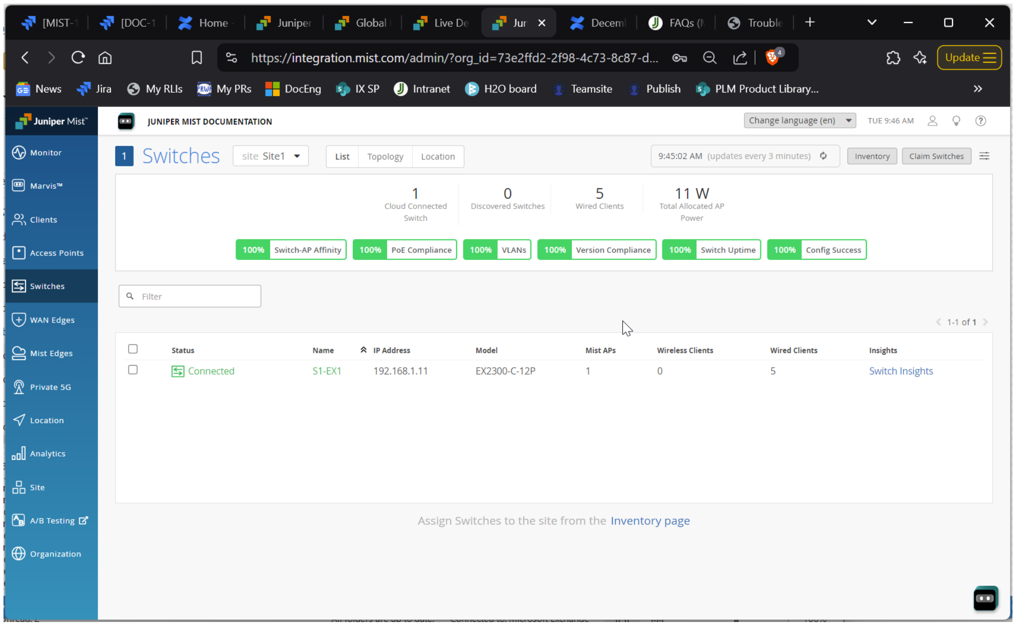

1 Cloud Connected Switch | 0 Discovered Switches | 5 Wired Clients | 11 W Total Allocated AP Power

100% Switch-AP Affinity | 100% PoE Compliance | 100% VLANs | 100% Version Compliance | 100% Switch Uptime | 100% Config Success

	Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	Insights
□	Connected	S1-EX1	192.168.1.11	EX2300-C-12P	1	0	5	Switch Insights

Assign Switches to the site from the Inventory page

Troubleshooting Steps

If the Juniper Mist™ portal shows a switch as disconnected when it is online and reachable locally, you can troubleshoot the issue. You need console access or SSH access to the switch to perform the troubleshooting steps listed in this topic.

To troubleshoot your switch:

1. Ensure that the Junos OS version running on the switch supports zero-touch provisioning (ZTP). For example, the EX2300 and EX3400 switches require Junos OS version 18.2R3-S2 or later. The EX4300 switch requires Junos OS 18.4R2-S2 or later. The EX4600 and EX4650 switches require Junos OS 20.4R3 or later.
2. Log in to the switch CLI and run `show interfaces terse`.

   ```
   user@switch> show interface terse
   Interface        Admin  Link   Proto       Local
   ge-0/0/0         up     up
   irb.0            up     up     inet        192.168.3.24/24
   me0              up     down
   me0.0            up     down   inet        192.168.3.24/24
   ...truncated...
   ```

   You should see the integrated routing and bridging (IRB) interface (irb.0) with an IP address. You might see multiple IRB interfaces, depending on the switch model (or in the case of a Virtual Chassis).

   At least one IRB interface needs to have a valid IP address. The switch can also connect using a management IP address, which you can see on the me0 interface. Ensure that either the irb0 or me0 interface has a valid IP address and has its Admin and Link states up.
3. Ensure that the switch can reach the gateway.
4. Use a ping test, as follows, to ensure that the switch can reach the Internet:

   ```
   user@switch> ping 8.8.8.8
   PING 8.8.8.8 (8.8.8.8): 56 data bytes
   64 bytes from 8.8.8.8: icmp_seq=0 ttl=117 time=22.996 ms
   64 bytes from 8.8.8.8: icmp_seq=1 ttl=117 time=24.747 ms
   64 bytes from 8.8.8.8: icmp_seq=2 ttl=117 time=16.528 ms

   --- 8.8.8.8 ping statistics ---
   3 packets transmitted, 3 packets received, 0% packet loss
   round-trip min/avg/max/stddev = 16.528/21.424/24.747/3.535 ms
   ```

5. Check if the switch can resolve `oc-term.mistsys.net` and `jma-term.xx.mistsys.net` by using a ping test. Sample ping tests are shown below:

```
user@switch> ping oc-term.mistsys.net
PING ab847c3d0fcd311e9b3ae02d80612151-659eb20beaaa3ea3.elb.us-west-1.amazonaws.com
(13.56.90.212): 56 data bytes
```

```
user@switch> ping jma-terminator-staging.mistsys.net
PING a8481a00030ad459aac15af07d5f2c5b-75855524.us-east-1.elb.amazonaws.com (3.210.247.53):
56 data bytes
^C
--- a8481a00030ad459aac15af07d5f2c5b-75855524.us-east-1.elb.amazonaws.com ping statistics
---
1 packets transmitted, 0 packets received, 100% packet loss
```

   If the switch is not resolving `oc-term.mistsys.net` or `jma-term.xx.mistsys.net`, make sure that the switch has a DNS server configured.

```
user@switch> show configuration | display set | grep name-server
set system name-server 202.56.230.2
set system name-server 202.56.230.7
set system name-server 8.8.8.8
```

   If the switch doesn't have a DNS server, configure the server as shown in the following example: `user@switch# set system name-server 8.8.8.8`

6. Ensure that the required firewall port (TCP port 2200 for oc-term.mistsys.net) is open.

```
user@switch> show system connections | grep 2200
tcp4 0 0 192.168.3.24.64647 13.56.90.212.2200 ESTABLISHED
```

   See Device-to-Cloud Addresses and Ports to determine which port to enable, depending on your cloud environment.

NOTE: The EX2300, EX3400, EX4100, EX4400, EX4650, EX5120 switches no loner need the port 2200. These switches connect to Mist cloud over HTTPS port 443. See also: "Troubleshooting Juniper CloudX" on page 332.

7. Check the system time on the switch to make sure the time is correct.

```
user@switch> show system uptime
fpc0:
--------------------------------------------------------------------------
Current time: 2020-09-01 21:49:05 UTC
Time Source: LOCAL CLOCK
System booted: 2020-08-27 06:57:04 UTC (5d 14:52 ago)
Protocols started: 2020-08-27 07:01:35 UTC (5d 14:47 ago)
Last configured: 2020-09-01 17:21:59 UTC (04:27:06 ago) by mist
9:49PM up 5 days, 14:52, 2 users, load averages: 0.79, 0.65, 0.58
```

   If the system time is not correct, configure it. For more information, see Configure Date and Time Locally.

8. Check `device-id` to make sure it is in the format `<org_id>.<mac_addr>`, as shown below:

```
user@switch# show system services outbound-ssh
traceoptions {
file outbound-ssh.log size 64k files 5;
flag all;
}
client mist {
device-id ca01ea19-afde-49a4-ad33-2d9902f14a7e.e8a2453e672e;
secret "$9$L7i7-wgoJUDkg49Ap0IRrevW-VYgoDHqWLGDkqQzRhcreWLX-Vs2XxGDHkPfn/Cp0IcSeMLxn/LxN-
ws5Qz6tuRhSv8Xrl87dVY2TzF/uOEcyKWLleUjikPfIEhSrvxNdbYgRhK8x7Vbk.mf5F9CuOBEtp0IcSMWoJZjmfFn/
CA05TIEhSeK4aJUjqP5Q9tu4an/CtOB7-dboJZUjHmfaJn/ApREevW8X-
YgoiqmxNb2gaUD69Cp1RSyKMLxCtORSrvM7-VboJDjqPTzNdmfzF/
9vW8LdbY2aZGisY4ZDif5z3690BylKWX7KvZUHkTQlKvW-VJGDiqmGU/
CtuEhKM87wYaJDkqfoaQFn6At1RhrM8xNd"; ## SECRET-DATA
keep-alive {
retry 3;
timeout 5;
}
services netconf;
oc-term.mistsys.net {
port 2200;
retry 1000;
timeout 60;
}
}
```

See outbound-ssh for more information.

You can also examine the log messages by using the command `show log messages`.

9. If you are adding the switch for the first time, do the following:

   - Delete the present Juniper Mist configuration from the switch using the delete command.
   - Onboard the switch again using the claim or adopt workflow.
   - Verify the system connection using the `show system connections | grep 2200` command. If the switch remains disconnected with the sessions stuck in FIN_WAIT state, but is able to reach the Internet and resolve DNS, check for any maximum transmission unit (MTU) issues.

10. To check for any MTU issues, initiate a ping test toward any public server (for example, 8.8.8.8).

    Another way to check for MTU issues is to review the uplink packet capture file from the switch. A failing transaction due to an MTU issue would look like the following example. The example shows that the packets with a size of 1514 are being retried.

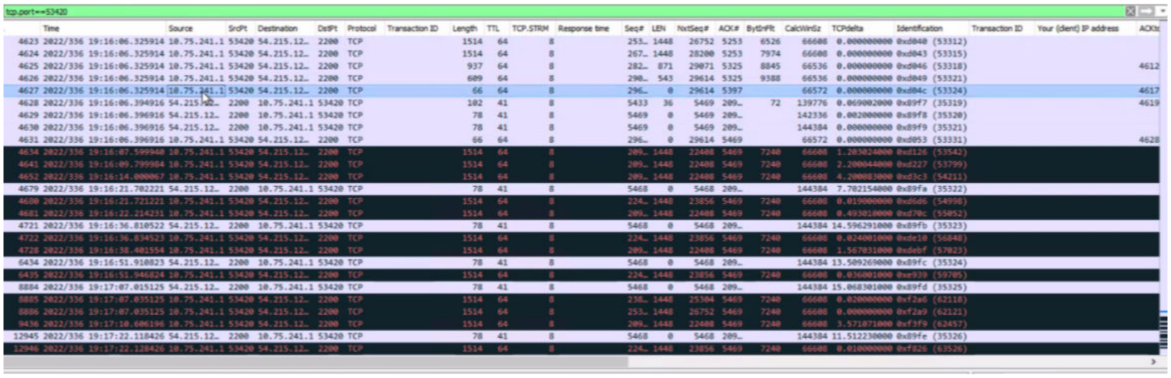

	Time	Source	SrcPt	Destination	DstPt	Protocol	Transaction ID	Length	TTL	TCP.STRM	Response time	Seq#	LEN	NxtSeq#	ACK#	BytInFlt	CalcWinSz	TCPdelta	Identification	Transaction ID	Your (client) IP address	ACKt
4623	2022/336 19:16:06.325914	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		253…	1448	26752	5253	6526	66608	0.000000000	0xd040 (53312)			
4624	2022/336 19:16:06.325914	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		267…	1448	28200	5253	7974	66608	0.000000000	0xd043 (53315)			
4625	2022/336 19:16:06.325914	10.75.241.1	53420	54.215.12…	2200	TCP		937	64	8		282…	871	29071	5325	8845	66536	0.000000000	0xd046 (53318)			4612
4626	2022/336 19:16:06.325914	10.75.241.1	53420	54.215.12…	2200	TCP		609	64	8		290…	543	29614	5325	9388	66536	0.000000000	0xd049 (53321)			
4627	2022/336 19:16:06.325914	10.75.241.1	53420	54.215.12…	2200	TCP		66	64	8		296…	0	29614	5397		66572	0.000000000	0xd04c (53324)			4617
4628	2022/336 19:16:06.394916	54.215.12…	2200	10.75.241.1	53420	TCP		102	41	8		5433	36	5469	209…	72	139776	0.069002000	0x89f7 (35319)			4619
4629	2022/336 19:16:06.396916	54.215.12…	2200	10.75.241.1	53420	TCP		78	41	8		5469	0	5469	209…		142336	0.002000000	0x89f8 (35320)			
4630	2022/336 19:16:06.396916	54.215.12…	2200	10.75.241.1	53420	TCP		78	41	8		5469	0	5469	209…		144384	0.000000000	0x89f9 (35321)			
4631	2022/336 19:16:06.396916	10.75.241.1	53420	54.215.12…	2200	TCP		66	64	8		296…	0	29614	5469		66572	0.000000000	0xd053 (53331)			4628
4634	2022/336 19:16:07.599940	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		209…	1448	22408	5469	7240	66608	1.203024000	0xd126 (53542)			
4641	2022/336 19:16:09.799984	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		209…	1448	22408	5469	7240	66608	2.200044000	0xd227 (53799)			
4652	2022/336 19:16:14.000067	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		209…	1448	22408	5469	7240	66608	4.200083000	0xd3c3 (54211)			
4679	2022/336 19:16:21.702221	54.215.12…	2200	10.75.241.1	53420	TCP		78	41	8		5468	0	5468	209…		144384	7.702154000	0x89fa (35322)			
4680	2022/336 19:16:21.721221	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		224…	1448	23856	5469	7240	66608	0.019000000	0xd6d6 (54998)			
4681	2022/336 19:16:22.214231	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		209…	1448	22408	5469	7240	66608	0.493010000	0xd70c (55052)			
4721	2022/336 19:16:36.810522	54.215.12…	2200	10.75.241.1	53420	TCP		78	41	8		5468	0	5468	209…		144384	14.596291000	0x89fb (35323)			
4722	2022/336 19:16:36.834523	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		224…	1448	23856	5469	7240	66608	0.024001000	0xde10 (56848)			
4728	2022/336 19:16:38.401554	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		209…	1448	22408	5469	7240	66608	1.567031000	0xdebf (57023)			
6434	2022/336 19:16:51.910823	54.215.12…	2200	10.75.241.1	53420	TCP		78	41	8		5468	0	5468	209…		144384	13.509269000	0x89fc (35324)			
6435	2022/336 19:16:51.946824	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		224…	1448	23856	5469	7240	66608	0.036001000	0xe939 (59705)			
8884	2022/336 19:17:07.015125	54.215.12…	2200	10.75.241.1	53420	TCP		78	41	8		5468	0	5468	209…		144384	15.068301000	0x89fd (35325)			
8885	2022/336 19:17:07.035125	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		238…	1448	25304	5469	7240	66608	0.020000000	0xf2a6 (62118)			
8886	2022/336 19:17:07.035125	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		253…	1448	26752	5469	7240	66608	0.000000000	0xf2a9 (62121)			
9436	2022/336 19:17:10.606196	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		209…	1448	22408	5469	7240	66608	3.571071000	0xf3f9 (62457)			
12945	2022/336 19:17:22.118426	54.215.12…	2200	10.75.241.1	53420	TCP		78	41	8		5468	0	5468	209…		144384	11.512230000	0x89fe (35326)			
12946	2022/336 19:17:22.128426	10.75.241.1	53420	54.215.12…	2200	TCP		1514	64	8		224…	1448	23856	5469	7240	66608	0.010000000	0xf826 (63526)			

    To troubleshoot this issue further, do a ping test from the switch. Use different ping sizes as shown in the following example:

```
user@switch> ping size 1450 8.8.8.8
PING 8.8.8.8 (8.8.8.8): 1450 data bytes
76 bytes from 8.8.8.8: icmp_seq=0 ttl=59 time=12.444 ms
— 8.8.8.8 ping statistics —
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 12.318/12.381/12.444/0.063 ms
```

As you can see below, the ping test with the size of 1480 has failed.

```
user@switch> ping size 1480 8.8.8.8
PING 8.8.8.8 (8.8.8.8): 1480 data bytes

– 8.8.8.8 ping statistics –
4 packets transmitted, 0 packets received, 100% packet loss
```

To resolve this issue, you can adjust the MTU on the uplink, based on the byte size at which packets are getting timed out.

11. Deactivate and then reactivate the outbound SSH, as shown below:

```
user@switch# deactivate system services outbound-ssh client mist
user@switch# activate system services outbound-ssh client mist
user@switch# commit
```

Watch the following video as well for more information on how to troubleshoot a switch:

Video: Wired Assurance Troubleshooting

Troubleshoot Disconnected Switches with Juniper Support

When a switch is disconnected from the Mist cloud cannot reconnect, you can quickly find out the disconnection reason using the switch CLI and share the details with the Juniper support team.

To check the disconnection reason, run the following CLI command locally on the switch: **op mist_debug.py cmd check-connectivity**.

NOTE: This command should be used only when the switch is disconnected from the Mist cloud. It retrieves the reason for the current disconnection state. This command does not work if the switch has never connected to the Mist cloud.

Example

```
user@switch> op mist_debug.py cmd check-connectivity
Starting Cloud Connection Troubleshooting
Management routing-instance is not configured
```

```
Checking if Switch IP is configured...
Switch IPs configured: ['178.35.0.45']
----------------------------------------
Checking if Gateway is configured...
Default gateway is configured: 178.35.0.1
----------------------------------------
Checking if Gateway is reachable...
Default gateway is not reachable
Failed
Troubleshooting ended with trouble code: GATEWAY_UNREACHABLE
```

Troubleshooting Juniper CloudX

SUMMARY

Follow these steps to investigate issues with communications between your switch and the Juniper Mist™ using CloudX.

For a list of platforms and Junos versions that support Juniper CloudX, refer to Juniper CloudX Overview.

For CloudX to work, you must ensure that the firewall port towards jma-terminator.xx.mist.com is open and SSL encryption is disabled on the firewall (for more information, refer to Juniper Mist Firewall Ports and IP Addresses for Firewall Configuration).

To check if a switch communicates with Mist cloud using CloudX:

1. Run the below CLI commands on the switch:

```
{master:0}
user@switch> show version | match mist
JUNOS Mist Agent [v1.0.2205-2]

{master:0}
user@switch> show system connections | grep 443
```

```
tcp4       0      0  192.168.2.52.62957
52.52.102.40.443                              ESTABLISHED
```

To verify CloudX through the Mist portal, you can use the steps below:

a. Log in to the Mist portal (manage.mist.com).

b. Click **Switches** > ***switch name*** to go to the switch detail page.

c. Click any port or a range of ports.

   If CloudX is running, the **Packet Capture** button is enabled; otherwise, the button is grayed out.

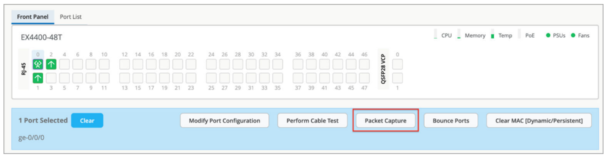

You can also check if CloudX is enabled on multiple switches by using the Mist portal.

To do that, click **Site** > **Switch Packet Captures** > **Add Switch**.

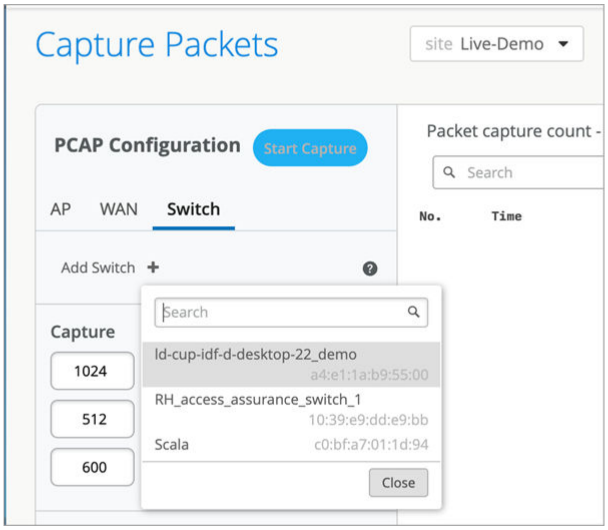

ld-cup-idf-d-desktop-22_demo a4:e1:1a:b9:55:00
RH_access_assurance_switch_1 10:39:e9:dd:e9:bb
Scala c0:bf:a7:01:1d:94

The switches listed here are all CloudX-enabled.

2. Verify that Mist Cloud Daemon (mcd) and Junos Mist Daemon (jmd) are running.

   mcd is responsible for enabling communication between the switch and the cloud. It maintains a secure WebSocket connection to the terminator in the cloud.

   jmd is used for:

   - Generating periodic statistics for the device.

   - Applying device configuration.

- Gathering device events.
- Initiating device functions (such as packet capture and software updates).
- Returning results from requested functions (such as files and streamed data).

To verify that jmd and mcd are running, use the following CLIs:

```
user@switch> set cli screen-width 400
user@switch> start shell
% ps aux | grep jmd
root   21408   0.0  0.4 1246080  32200  -  S    Fri23      15:17.51 /var/run/scripts/jet/jmd -
mcd-socket /var/run/mist_mcd.ipc
mist    3706   0.0  0.0   11136   2516  0  S+   07:14       0:00.00 grep jmd
%
%
% ps aux | grep mcd
root   21319   0.0  0.3 1242924  22256  -  I    Fri23       8:18.00 /var/run/scripts/jet/mcd
root   21408   0.0  0.4 1246080  32200  -  S    Fri23      15:17.53 /var/run/scripts/jet/jmd -
mcd-socket /var/run/mist_mcd.ipc
mist    3708   0.0  0.0   11136   2516  0  S+   07:14       0:00.00 grep mcd
%
```

3. Check the jmd and mcd logs for any errors by using the CLI commands below. Typically, jmd logs shows issues related to configuration or stats. The mcd logs report issues related to the connectivity between the switch and the cloud.

```
user@switch> show log jmd.log | last 10
[jmd] 2024/11/04 07:12:02 collector.go:850: total stats collection time = 10s
[jmd] 2024/11/04 07:12:02 app_states.go:355: app sending stats to mist cloud (26171 bytes)
[jmd] 2024/11/04 07:12:02 app_states.go:360: successfully sent ipc stats:
[jmd] 2024/11/04 07:12:02 app.go:282: processing app state "STEADY"
[jmd] 2024/11/04 07:12:12 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:22 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:32 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:42 app.go:339: sending ipc keep-alive
```

```
[jmd] 2024/11/04 07:12:52 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:52 collector.go:417: collecting periodic stats, interval 60
```

```
user@switch> show log mcd.log | last 10
[mcd] 2024/11/04 07:09:31 app.go:967: successfully sent msg to cloud: ep-telemetry
[mcd] 2024/11/04 07:10:02 ipc_server.go:414: rx ipc request: send cloud telemetry
[mcd] 2024/11/04 07:10:02 ipc_server.go:447: forwarding ipc telemetry to "junos-stats-"
(26167 bytes)
[mcd] 2024/11/04 07:11:02 ipc_server.go:414: rx ipc request: send cloud telemetry
[mcd] 2024/11/04 07:11:02 ipc_server.go:447: forwarding ipc telemetry to "junos-stats-"
(26171 bytes)
[mcd] 2024/11/04 07:12:01 app.go:967: successfully sent msg to cloud: ep-telemetry
[mcd] 2024/11/04 07:12:02 ipc_server.go:414: rx ipc request: send cloud telemetry
[mcd] 2024/11/04 07:12:02 ipc_server.go:447: forwarding ipc telemetry to "junos-stats-"
(26171 bytes)
[mcd] 2024/11/04 07:13:02 ipc_server.go:414: rx ipc request: send cloud telemetry
[mcd] 2024/11/04 07:13:02 ipc_server.go:447: forwarding ipc telemetry to "junos-stats-"
(26171 bytes)
```

4. If jmd or mcd is not running for some reason, try restarting it, as shown in the sample below.

```
{master:0}
user@switch> request extension-service restart-daemonize-app mcd
Extension-service application 'mcd' with pid: 92502 exited with return: -1
Extension-service application restarted successfully
```

> **NOTE**: If you are using a 22.4xx Junos version, use the command `request extension-service daemonize-restart mcd`.

5. If the switch is not connecting to the cloud, check its reachability by using a ping and curl test. These tests will help you check if the required firewall ports are allowed.

   The cloud endpoints are not set up to respond to ping tests; however, running a ping test (as shown below) will ensure that DNS resolves FQDN.

```
user@switch> ping jma-terminator-cloud_env.mist.com
```

Example:

```
user@switch> ping jma-terminator-ac2.mist.com
PING a8481a00030ad459aac15af07d5f2c5b-75855524.us-east-1.elb.amazonaws.com (3.210.247.53): 56
data bytes
^C
--- a8481a00030ad459aac15af07d5f2c5b-75855524.us-east-1.elb.amazonaws.com ping statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
```

Here is a sample curl test. Replace the URL in the below sample with your designated URL.

```
user@switch> start shell
% curl -k https://jma-terminator.mistsys.net/test
Welcome to MIST
%
```

A valid response from the curl test proves that the jma-terminator in the Mist cloud is reachable. A lack of response or receipt of an error will indicate that the path between the switch and the cloud is blocking these ports, likely because of the firewall. The URLs used in the test are the same as those in firewall ports and differ between cloud instances.

RELATED DOCUMENTATION

Juniper CloudX Overview

Cloud-Ready LED Blink Patterns

SUMMARY

Use this table to understand what the flashing CLD LED is telling you about your switch's ability to connect to cloud services.

IN THIS SECTION

The following table tells what the different blinking **CLD** LED patterns mean and what you can do to address it. In addition to observing the physical switch, you can also assess the status from the Junos CLI by issuing the `show chassis led` command.

Note that for Virtual Chassis (VC) deployments managed from the Mist cloud, the CLD LED reflect the state of the primary, except when a software download is in progress (in which case all members of the VC will show OS upgrade blink pattern and color).

Table 26: Cloud LED Blink Patterns

CLD LEDs	Blink Pattern	Meaning
●	solid green	The ZTP process is complete.
○	solid white	Connected to Mist cloud.
● ● ●	3 yellow	No IP Address. The DHCP server is not configured or could not be reached. Junos did not receive a DHCP lease or IP address.
● ● ● ●	4 yellow	No default gateway. Either the address was not received or it is not configured on the device,
● ● ● ● ●	5 yellow	The default gateway could not be reached. No ARP from the default gateway.
● ● ● ● ● ●	6 yellow	No DNS server(s) found in the static configuration, or in the DHCP lease.

Table 26: Cloud LED Blink Patterns *(Continued)*

CLD LEDs	Blink Pattern	Meaning
● ● ● ● ● ● ●	7 yellow	No response from the DNS server. The switch received an IP address for the DNS server via DHCP, but it cannot not reach the Mist cloud.
● ● ● ● ● ● ● ●	9 yellow	The Mist agent cannot reach the Mist cloud.
● ○ ● ●	1 yellow, pause, 2 yellow	Could not connect to the redirect server, most likely due to a firewall blocking TCP port 443, TCP port 2200. See also Ports to Open in Your Firewall.
● ○ ● ● ● ●	1 yellow, pause, 4 yellow	Invalid configuration on the redirect server (PHC). This device received a 500 or 404 error from the redirect server at **redirect.juniper.net**.
● ○ ● ● ● ● ●	1 yellow, pause, 5 yellow	Incorrect time on the switch. During ZTP, the phone home client (PHC) received a certificate with the wrong time. ZTP could not continue.

Table 26: Cloud LED Blink Patterns *(Continued)*

CLD LEDs	Blink Pattern	Meaning
● ○ ● ● ● ● ● ●	1 yellow, pause, 6 yellow	Cloud unreachable. During ZTP, the PHC could not reach the cloud.

Cloud-Ready Connection Process

SUMMARY

Learn about the process that your switches use to connect to the services in the Juniper Mist™ cloud.

Juniper EX Series switches are cloud-ready devices, which means they are Day-0 capable of connecting to the Juniper Mist cloud. When running a supported version of Junos, these switches can also automatically establish a connection to Juniper Mist cloud services, where they can then be on-boarded (Day 1), managed (Day 2), and monitored (Day2+) from the Mist portal.

As part of zero-touch-provisioning (ZTP), a secure TCP connection uses pre-shared keys on both the device and cloud to establish a connection. Figure 1 provides a break-down of the ZTP process.

On the front of EX4100, EX4100-F, and EX4400 switches there is a CLD interface LED that you can use to monitor the ZTP progress. The blink pattern of the LED can help troubleshoot any Day 0 connection issues.

For switches already in the cloud, you can view the connection status from the Switches page of the Mist portal.

Figure 30: Stages in the ZTP Process for Cloud Ready Switches

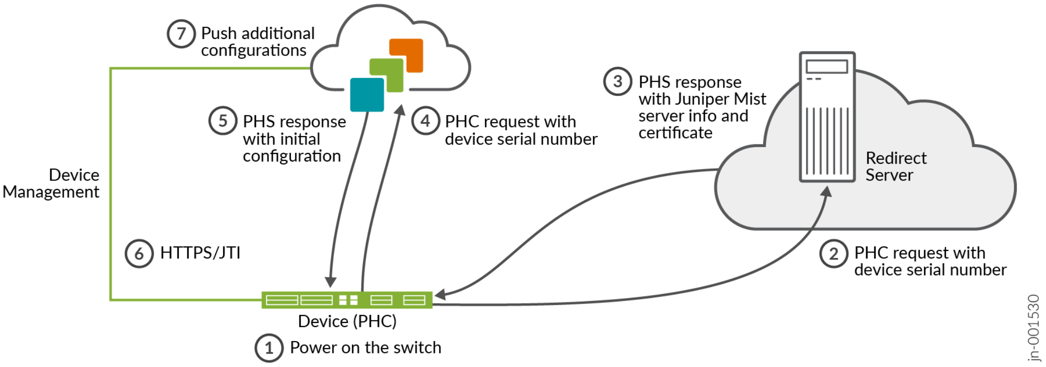
7 Push additional configurations
5 PHS response with initial configuration
4 PHC request with device serial number
3 PHS response with Juniper Mist server info and certificate
Redirect Server
Device Management
6 HTTPS/JTI
2 PHC request with device serial number
Device (PHC)
1 Power on the switch
jn-001530

The first time a cloud-ready switch is powered on, an on-board phone-home client (PHC) connects to a redirect server, which then redirects it to a phone home server (PHS) where the switch can get the latest Junos configuration. You can also have the switch connect to a DHCP server that supports ZTP and run the ZTP process from there.

NOTE:

Switches running outbound SSH are managed using SSH/NETCONF sessions. Switches enabled with CloudX use HTTPS and JTI for management, which is the default behavior.

RELATED DOCUMENTATION

Cloud-Ready LED Blink Patterns | 336

RELATED DOCUMENTATION

Cloud-Ready Connection Process | 339

Troubleshoot with Marvis

SUMMARY

Watch this video to learn about the troubleshooting help that's available from the Marvis® Virtual Network Assistant.

Marvis® Virtual Network Assistant is an AI-driven, interactive virtual network assistant that streamlines network operations, simplifies troubleshooting, and provides an enhanced user experience. With real-time network visibility, Marvis provides a comprehensive view of your network from an organizational level to a client level with detailed insights. Marvis leverages the Mist AI to identify issues proactively and provide recommendations to fix issues.

To use Marvis for switches, you must have the Marvis for Wired subscription in association with the Wired Assurance base license.

Marvis can automatically fix issues (self-driving mode) or recommend actions that require user intervention (driver-assist mode). The Marvis Actions page lists the high-impact network issues that Marvis detects. Marvis Actions also displays the recommended actions for your organization's network.

For more information about Marvis actions for switches, see Marvis Actions for Switches.

Video: Marvis Actions for Switches

FAQs (Mist Wired)

SUMMARY

Get answers to common questions about missing VLANs, default configuration options, port configuration, upgrades, and more,

IN THIS SECTION

What does the Inactive wired VLANs warning on the Mist dashboard mean?

When your APs do not detect incoming traffic from a particular VLAN that is used in either an AP or a WLAN configuration, Mist suspects that this VLAN is not configured on the switch port where the APs are connected. The Inactive wired VLANs warning appears on the AP list page to indicate this issue, and an icon is displayed next to the APs experiencing the inactive wired VLAN issue.

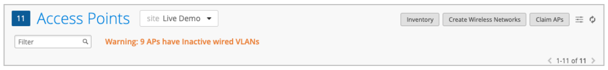

Warning: 9 APs have Inactive wired VLANs

How to check which VLAN is missing on the switch port?

To find out which VLANs are missing on the switch port:

1. Go to the Marvis Actions page (**Marvis** > **Actions**).

2. From the actions tree, select **Switch** > **Missing VLAN** to see the VLANs that are missing.

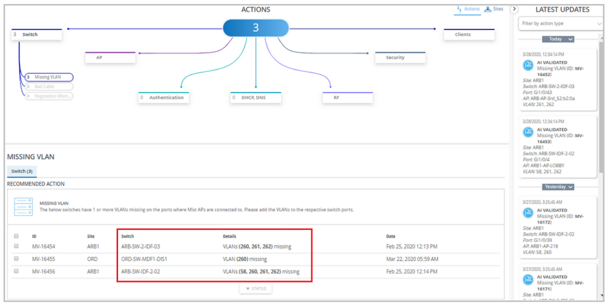

MISSING VLAN

Switch (3)

RECOMMENDED ACTION

MISSING VLAN

The below switches have 1 or more VLANs missing on the ports where Mist APs are connected to. Please add the VLANs to the respective switch ports.

ID	Site	Switch	Details	Date
MV-16454	ARB1	ARB-SW-2-IDF-03	VLANs (260, 261, 262) missing	Feb 25, 2020 12:13 PM
MV-16455	ORD	ORD-SW-MDF1-DIS1	VLAN (260) missing	Mar 22, 2020 05:59 AM
MV-16456	ARB1	ARB-SW-IDF-2-02	VLANs (58, 260, 261, 262) missing	Feb 25, 2020 12:14 PM

STATUS

3/28/2020, 12:34:14 PM
AI VALIDATED
Missing VLAN (ID: **MV-16452**)
Site: ARB1
Switch: ARB-SW-2-IDF-03
Port: Gi1/0/43
AP: ARB-AP-3rd_52:b2:0a
VLAN: 261, 262

3/28/2020, 12:34:14 PM
AI VALIDATED
Missing VLAN (ID: **MV-16453**)
Site: ARB1
Switch: ARB-SW-IDF-2-02
Port: Gi1/0/4
AP: ARB1-AP-LOBBY
VLAN: 58, 261, 262

3/27/2020, 3:25:45 AM
AI VALIDATED
Missing VLAN (ID: **MV-16172**)
Site: ARB1
Switch: ARB-SW-IDF-2-02
Port: Gi1/0/39
AP: ARB1-AP-218
VLAN: 58, 260

3/27/2020, 3:25:45 AM
AI VALIDATED
Missing VLAN (ID: **MV-16171**)
Site: ARB1

How to verify if Marvis is detecting the correct case of missing VLANs?

To verify whether the Marvis AI is detecting the correct case of missing VLANs, do a packet capture or port mirroring on the switch port to which the AP is connected, and use the Wireshark tool to analyze the traffic. You can also use the VLAN filter to verify if any traffic is coming from that VLAN. See Dynamic and Manual Packet Captures for more help on setting up Wireshark.

How to fix the missing VLAN error?

Once you have identified the VLAN that is missing from the switch port but is being used in your AP or WLAN configuration, you can configure that VLAN on your switch. After the VLAN is correctly configured on your switch and the AP starts detecting traffic from it, Mist takes some time to verify the fix and ensure that the issue is resolved. After that, the warning disappears automatically.

NOTE: If you see the warning even after fixing all the VLANs on your switch ports, open a support ticket for assistance. For more information, see Create a Support Ticket.

In an EX Series switch-based Virtual Chassis (VC) system, how do I convert the DAC-attached VC port to a trunk port?

You can do this by converting the VC port on the Virtual Chassis to network port by running the following CLI command locally on the switch:

```
request virtual-chassis mode network-port
```

This behavior applies to the EX Series switches that come with a default VC port. This is a one-time step.

Is Spanning Tree Protocol (STP) enabled on EX Series switches by default?

STP is not enabled on EX Series switches by default. However, if the switch is managed by Mist, this protocol is enabled on the switch automatically. If users do not want STP on a port, they can set the port as a Rapid Spanning Tree Protocol (RSTP) edge port by using a port profile. This setting ensures that the port is treated as an edge port and guards against the reception of Bridge Protocol Data Units (BPDUs).

I created a port configuration and applied that to a range of ports on a switch. Now, I want to select a single port from that range and configure an individual port description on it. How do I do that?

To do that, follow the steps below:

1. Go to the switch details page (**Switches** > *Switch Name*).
2. Select the port to be updated from the Front Panel section.
3. Click the Modify Port Configuration button to open the Port Configuration tile.
4. Specify a port description in the Description field.
5. Click the check mark (✓) on the top right of the Port Configuration tile to save the configuration.

NOTE: Saving this port configuration will remove the overlapping port IDs from the existing configurations (applied to the range of ports) and create a new port configuration for the selected port.

My switch lost connectivity during the onboarding process to Mist because of a phone-home client issue. How do I troubleshoot this issue?

You can try the below steps:

Check if the phone home process is running. To do that, log in to the switch CLI and run the command below:

```
show system processes extensive | match phone-home
```

If the phone-home process is not running, configure the phone-home server using the CLI commands below:

```
set system phone-home server https://redirect.juniper.net

set system phone-home rfc-compliant

commit
```

If the phone-home configuration is available on the device but the issue is persisting, restart the home client, using the CLI command below:

```
restart phone-home-client
```

If you don't want to carry out the above steps, you can consider zeroising the switch using the CLI command `request system zeroize`. This command will restart all the services on the switch, including the phone-home process.

If you want to collect logs from the phone-home process, run the CLI command `show log phc.log | last 100` on the switch.

My switch lost connectivity while being onboarded to the Mist cloud because of a native VLAN mismatch between the switch and the uplink device. How do I resolve this issue?

To resolve this issue, configure the uplink interface on the new switch that is being onboarded with the same native VLAN as the connected uplink device.

You can configure the native VLAN on the new switch uplink port through the template or directly from the switch details page. On the template (**Organization** > **Switch Templates**), you can configure the native VLAN from the IP Config tab on the Rules tile in the Select Switches Configuration section.

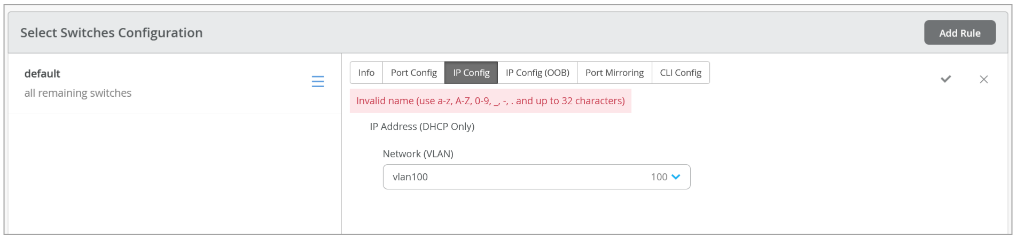

Invalid name (use a-z, A-Z, 0-9, _, -, . and up to 32 characters)

IP Address (DHCP Only)

Network (VLAN): vlan100 100

In the switch details page (**Switches** > **Switch Name**), you can configure the native VLAN at the device level from the IP Configuration tile in the Device section.

While being onboarded to the Mist cloud, my switch lost connectivity because the port on the uplink device does not have a native (untagged) VLAN with DHCP, DNS and Internet access. How do I resolve this issue?

To resolve this issue, try the below steps:

1. Create a configuration template for the new switch with uplink port profile configured with all VLANs tagged.

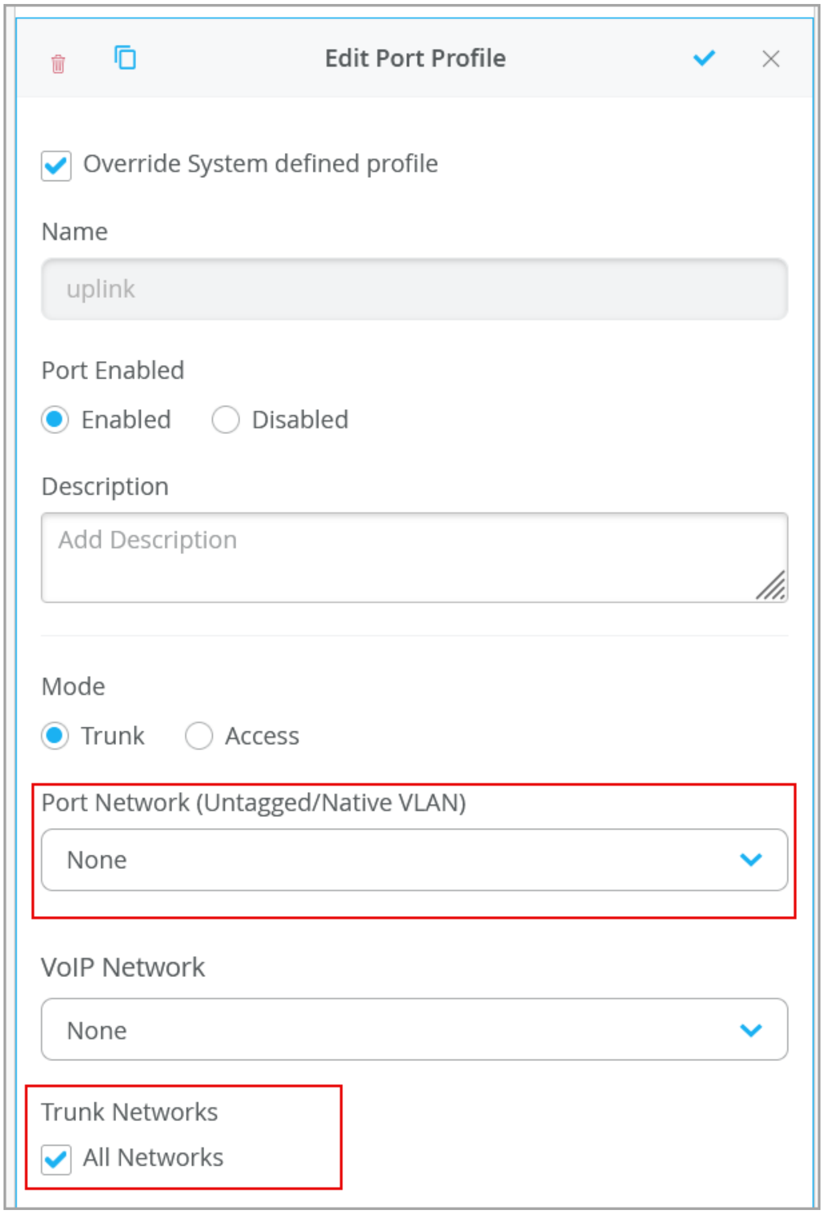

Edit Port Profile

☑ Override System defined profile

Name

uplink

Port Enabled

● Enabled ○ Disabled

Description

Add Description

Mode

● Trunk ○ Access

Port Network (Untagged/Native VLAN)

None

VoIP Network

None

Trunk Networks

☑ All Networks

2. Stage the switch by connecting its OOB-MGMT port to a network with DHCP, DNS, and internet access to start the ZTP process.

3. Once the configuration, which has uplink ports configured with all VLANs tagged, is pushed to the switch, disconnect the switch from OOB-MGMT and connect it to the existing switch via the revenue uplink ports for in-band management.

Upgrade Required: Why am I seeing an upgrade notification for my EX4400 switches, and how do I perform Junos or BIOS upgrade from the Mist dashboard?

For Juniper EX Series switches that are managed by Mist, you can identify and upgrade the Junos version and/or switch BIOS to the current version from the Mist dashboard. To alert you to a version mismatch and to recent releases, a notification appears at both the switch list, and on the configuration page of individual switch. When more than one switch is affected, you can select all to perform a simultaneous upgrade.

Figure 31: Upgrade Notification

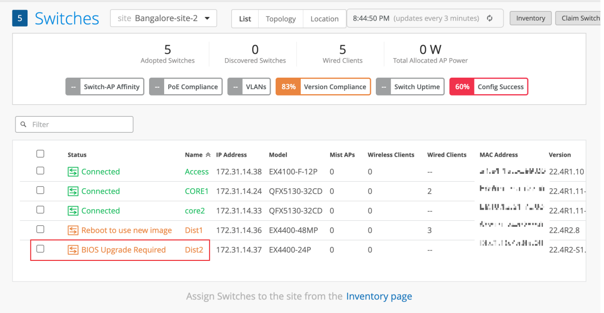

5 Switches site Bangalore-site-2 | List | Topology | Location | 8:44:50 PM (updates every 3 minutes) | Inventory | Claim Switch

5 Adopted Switches | 0 Discovered Switches | 5 Wired Clients | 0 W Total Allocated AP Power

-- Switch-AP Affinity | -- PoE Compliance | -- VLANs | 83% Version Compliance | -- Switch Uptime | 60% Config Success

Filter

□	Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	MAC Address	Version
□	Connected	Access	172.31.14.38	EX4100-F-12P	0	0	--	[redacted]	22.4R1.10
□	Connected	CORE1	172.31.14.24	QFX5130-32CD	0	0	2	[redacted]	22.4R1.11-
□	Connected	core2	172.31.14.33	QFX5130-32CD	0	0	--	[redacted]	22.4R1.11-
□	Reboot to use new image	Dist1	172.31.14.36	EX4400-48MP	0	0	3	[redacted]	22.4R2.8
□	BIOS Upgrade Required	Dist2	172.31.14.37	EX4400-24P	0	0	--	[redacted]	22.4R2-S1.

Assign Switches to the site from the Inventory page

When you see an upgrade notification, we recommend that you perform the update as soon as convenient to prevent "silent," or unexpected, device reboots that can result in the case of new Mist features not being supported by a previous version of the switch firmware. Note that for BIOS upgrades and some Junos upgrades, the switch will be taken offline to reboot or restart a given line card.

The **Upgrade BIOS** option appears:

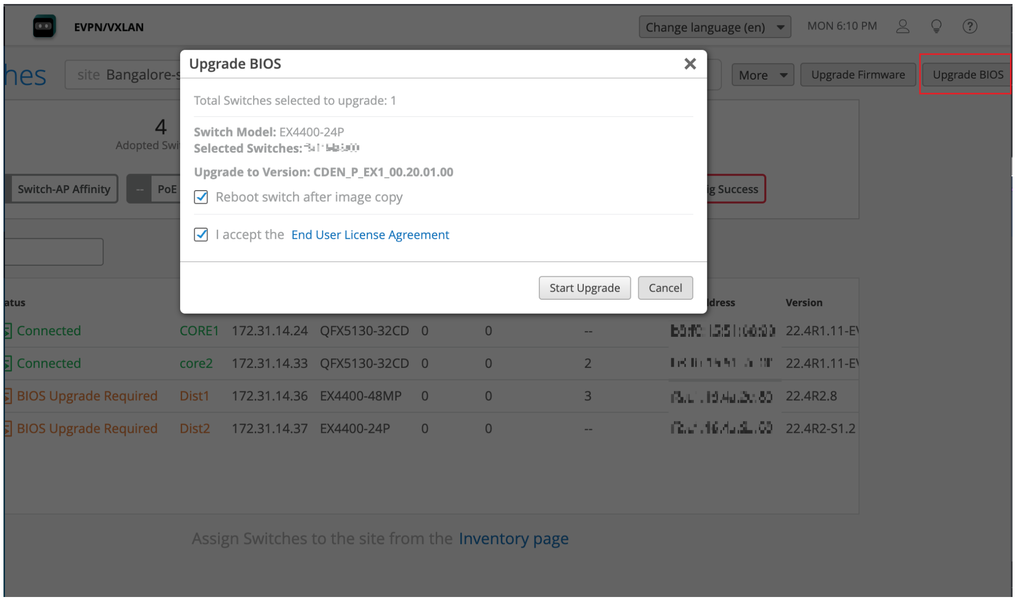

Upgrade BIOS ✕

Total Switches selected to upgrade: 1

Switch Model: EX4400-24P
Selected Switches: [redacted]

Upgrade to Version: CDEN_P_EX1_00.20.01.00

☑ Reboot switch after image copy

☑ I accept the End User License Agreement

Start Upgrade | Cancel

atus							ldress	Version
Connected	CORE1	172.31.14.24	QFX5130-32CD	0	0	--	[redacted]	22.4R1.11-E
Connected	core2	172.31.14.33	QFX5130-32CD	0	0	2	[redacted]	22.4R1.11-E
BIOS Upgrade Required	Dist1	172.31.14.36	EX4400-48MP	0	0	3	[redacted]	22.4R2.8
BIOS Upgrade Required	Dist2	172.31.14.37	EX4400-24P	0	0	--	[redacted]	22.4R2-S1.2

Assign Switches to the site from the Inventory page

For switches configured for virtual chassis (VC), all members are automatically upgraded together; you do not need to perform separate BIOS upgrades. Likewise, in the case where different FPCs are running different BIOS versions, only those FPCs that need the newer BIOS will be upgraded.

Tip: You can find the switch BIOS by opening a remote shell (**Utilities** > **Remote Shell** at the top of the switch configuration page) and running the following Junos command from the CLI:

```
show chassis firmware
```

The exact output will depend on the switch and hardware configuration, as shown in the following sample:

```
{master:0} user@device> show chassis firmware
                                    Part       Type       Version
           FPC 0                    loader     FreeBSD EFI loader 2.1
                                    BIOS       CDEN_P_EX1_00.15.01.00
                                    System CPLD 0.f
                                    CPU CPLD   1.0
           FPC 1                    loader     FreeBSD EFI loader 2.1
                                    BIOS       CDEN_P_EX1_00.20.01.00
                                    System CPLD 0.12
                                    CPU CPLD   1.0
           {master:0} user@device>
```

The following tables provide upgrade recommendations with regards to a known, and since September, 2023 fixed, issue with silent reboots on the EX4400 platform. If you have a device that is effected, upgrade the BIOS and/or Junos version as instructed in the following tables.

Table 27: Junos Upgrade Warning

Junos Version	Next Step	Comment
Junos 22.2R3-S1.11 and below.	Upgrade to Junos 22.2R3-S2.	Schedule a maintenance window and perform the upgrade. Multiple EX4400s can be upgraded together. Go to **Utilities > Upgrade Firmware** on the Switch detail page to perform the upgrade.
Junos 22.3	Upgrade to Junos 22.3R2-S2.	
Junos 22.4 and below 22.4R2-S1	Upgrade to Junos 22.4R2-S1.	

Table 28: BIOS Upgrade Warning

BIOS Upgrade Required	Next Step	Comment
If a warning for "BIOS Upgrade Required" is displayed on the Mist UI...	Upgrade the affected switches to the recommended BIOS version. If the switch is configured for VC, and is operating in HiGiG mode `network-port` or `HGoE mode` you will need to manually re-enable that virtual-chassis mode on the device and reboot. No such action is needed if the switch is running the default VC mode, `HiGiG mode`. To see which mode VC is running in, open a remote shell to the switch and run the following commands:`show virtual-chassis mode` and `show virtual-chassis vc-port`	Upgrading the BIOS upgrade requires the switch to reboot. Only those versions needing the upgrade will be upgraded, for example in a virtual chassis. On the Switch Detail page, go to **Utilities > Upgrade BIOS**. Select the EX4400 switches with the warnings and click **Upgrade BIOS**. For additional details, see TSB71527.

Table 29: BIOS and Junos Upgrade Warning

Both Warnings	Next Step	Comment
If a warning for BIOS as well as Junos Upgrade Required is displayed on the Mist UI...	We recommend that you perform both upgrades and then reboot once to minimize downtime.	On the Switch List page, select the affected switches and click **Upgrade Switches (without Reboot)** and **Upgrade BIOS (with Reboot)** to upgrade both Junos and BIOS in a single reboot.

How many sessions are recommended per single SSH connection for switches?

We recommend that you configure the default value which is 10. Range: 1 through 65535. See max-sessions-per-connection.

How many SSH connections are recommended for switches?

We recommend that you configure the default value which is 10. Range: 1 through 250. See connection-limit.

RELATED DOCUMENTATION

Manage or Update Configuration Settings | **149**

7

CHAPTER

Appendix

IN THIS CHAPTER

Deploy and Manage EX Series Switches at a Branch

SUMMARY

Get started with branch deployments by using this Juniper Validated Design document.

You can use Juniper Networks® EX Series Switches at branch locations as traditional standalone switches or as a Virtual Chassis combining up to ten switches and managing them as a single device. For higher scale deployments at the branch, you can use topologies that contains distribution switches between access switches and WAN devices.

The following document takes you through the workflow involved in deploying and managing EX Series Switches at the branch using the Juniper Mist™ cloud.

Distributed Branch EX Series—Juniper Validated Design (JVD).