Juniper Mist Wired Assurance Configuration Guide

Published
2024-03-03

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 Wired Service Levels

7 Troubleshooting

About This Guide

Hey there! If you're looking to configure Juniper switches using the Mist portal, you've come to the right place. The Mist portal offers essential features for switch configuration and management through the Juniper Mist Assurance cloud service.

Just a heads up, if you want to configure switches through the portal, make sure you have a Mist Super User role assigned to your account. Having this role will give you the necessary permissions and access to perform switch configurations within the portal. Happy configuring!

1

CHAPTER

Get Started

Juniper Mist Wired Assurance Overview

Juniper Mist™ Wired Assurance is an AI-driven cloud service that brings some awesome benefits, such as 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

Switch Configuration Overview (Mist) | 17

Hardware for Your Wired Network

SUMMARY

Read this topic to learn about the various hardware options and get started installing and onboarding your devices.

Juniper provides a wide range of hardware to support your wired networking needs. Use these links to find datasheets, quick start guides, and hardware guides.

- EX Series Switches
- QFX Series Switches

Switch Administrator Role Requirements

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	✓	×	×	×	×

(Continued)

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

(Continued)

Privileges	Super User	Network Admin (All Sites Access)	Network Admin (Site Group or Specific Sites Access)	Helpdesk	Observer
Enable/disable switch configuration management	✓	✓	✓	×	×
Send the switch logs to the Mist cloud	✓	✓	✓	×	×
View the Inventory page	✓	✓	×	×	✓
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	✓	✓	×	×	×

(Continued)

Privileges	Super User	Network Admin (All Sites Access)	Network Admin (Site Group or Specific Sites Access)	Helpdesk	Observer
Access to wired SLE, wired clients, the wired insights switch, or wired insights clients	✓	✓	✓	✓	✓

Deploy Your Wired Network

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
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 3
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 18

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

Category	Task	More Information
Configure Switches	Get started configuring your switches. For large-scale deployments, we recommend using switch configuration templates. Instead of configuring each switch individually, you can use a configuration template to set up and streamline configurations across multiple sites.	"Configure Switches" on page 22

Explore Juniper Mist Features

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 76
  - "Switch Details" on page 77
  - "Switch Utilities" on page 84
- Wired Service Level Expectations (SLEs)—Use the SLE dashboards to assess the network's user experience and resolve any issues proactively. See "Wired Service-Level Expectations (SLEs)" on page 154 .
- Port Profiles—Port profiles provide a convenient way to manually or automatically provision switch interfaces. See "Port Profiles Overview" on page 9 .
- 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 "Which Campus Fabric Topology to Choose" on page 120 .

- 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 "Group-Based Policy Configuration Overview (Mist)" on page 15 .

- 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" on page 87 .

Port Profiles Overview

IN THIS SECTION

Port profiles provide a convenient way to manually or automatically provision switch interfaces. Mist supports the following two types of port profiles based on how a profile is assigned to a port:

- 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).

Static Port Profiles

The static port profile assignment involves two steps - configuring a port profile and assigning it manually to a specific switch port. You can configure port profiles from the Port Profiles tile on the switch template or the switch details page. You can manually assign the profile to a 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.

Video: Port Profiles

Dynamic Port Profiles

Dynamic port profiles enable you to configure rules for dynamically assigning port profiles to an interface. 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.

Dynamic port configuration involves two steps:

1. Set up rules for dynamically 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, it assigns the 'AP' profile to the connected device.

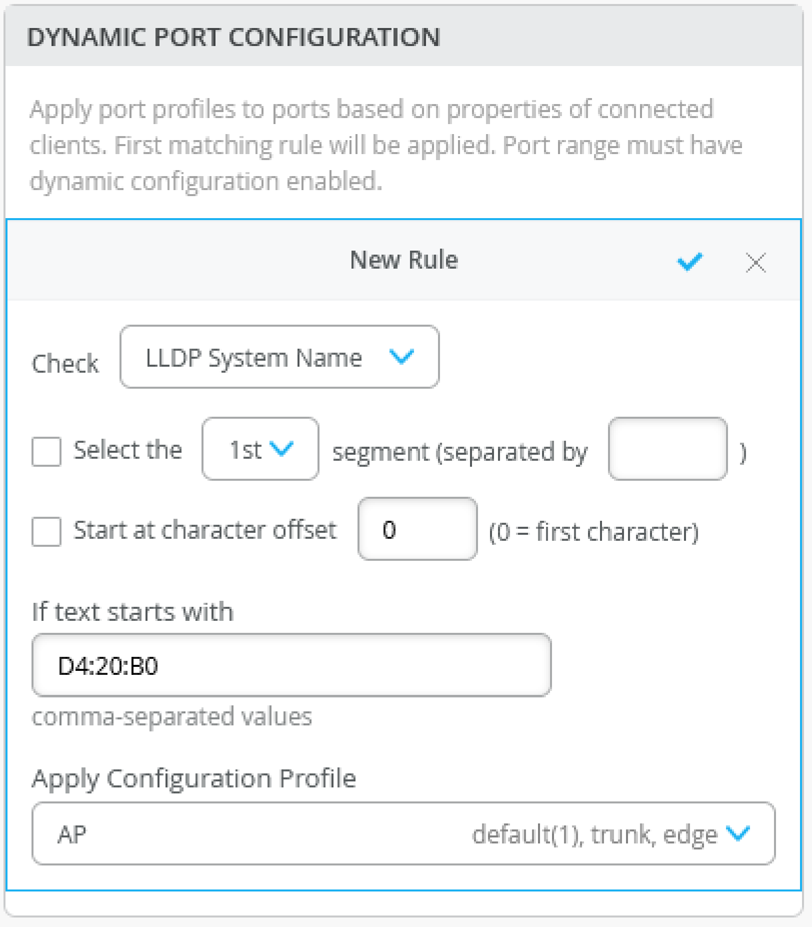

DYNAMIC PORT CONFIGURATION

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

comma-separated values

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

For more information, see the Dynamic Port Configuration step in "Create a Switch Configuration Template" on page 22 .

2. 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. You can also do this at the switch level, from the Port Configuration section on the switch details page.

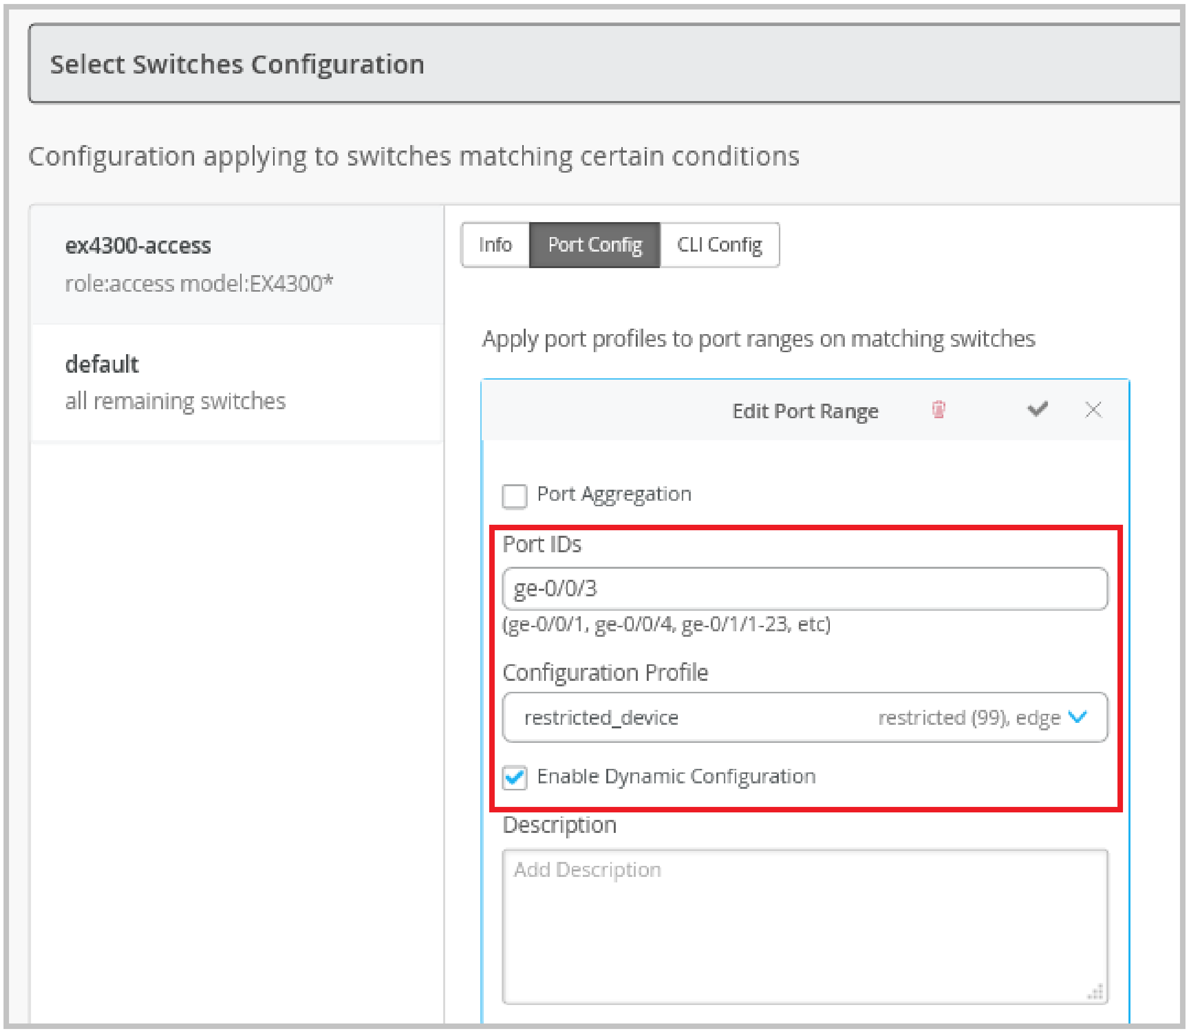

Select Switches Configuration

Configuration applying to switches matching certain conditions

ex4300-access
role:access model:EX4300*

default
all remaining switches

Apply port profiles to port ranges on matching switches

Edit Port Range

□ Port Aggregation

Port IDs: ge-0/0/3
(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Configuration Profile: restricted_device restricted (99), edge

☑ Enable Dynamic 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 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.

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.

See "Create a Switch Configuration Template" on page 22 for more information on how to configure port profiles.

Video: Dynamic Port Profiles (for Colorless Ports)

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.

Here is a sample port configuration for a Juniper EX Series switch. This configuration assumes the existence of a dedicated management VLAN, a staff VLAN, and a guest VLAN.

```
interfaces {
    ge-0/0/0 {
        native-vlan-id 100;
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members [ management staff guest ];
                }
            }
        }
    }
}

vlans {
    guest {
        vlan-id 667;
    }
    staff {
        vlan-id 200;
    }
    management {
        vlan-id 100;
        l3-interface irb.100;
    }
}
```

The following example shows how to set an IP address on the management VLAN of a switch (10.10.100.50/24) to be accessible from other networks (gateway of 10.10.100.1).

```

interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [ management staff guest ];
                }
                native-vlan-id 100;
            }
        }
    }
    vlan {
        unit 100 {
            family inet {
                address 10.10.100.50/24;
            }
        }
    }
}

routing-options {
    static {
        route 0.0.0.0/0 next-hop 10.10.100.1;
    }
}

vlans {
    guest {
        vlan-id 667;
    }
    staff {
        vlan-id 200;
    }
    management {
        vlan-id 100;
        l3-interface vlan.100;
    }

}
```

NOTE: For Juniper EX switches, we recommend that you include your switch's management address in the LLDP configuration:

In this example, the VLAN 100 is used for management, and the same is advertised over LLDP.

The following sample configuration is shown in set mode.

```
set interfaces irb unit 400 family inet address 10.33.1.110/24
set routing-options static route 0.0.0.0/0 next-hop 10.33.1.1
set routing-options static route 0.0.0.0/0 no-resolve
set protocols lldp management-address 10.33.1.110
set protocols lldp port-id-subtype interface-name
set protocols lldp interface all
set protocols lldp-med interface all
```

Group-Based Policy Configuration Overview (Mist)

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. 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.

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:

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:

See also: Microsegmentation with GBP Using Mist Wired Assurance.

2

CHAPTER

Switch Configuration

Switch Configuration Overview (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 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 CLI commands in the `set` format to configure additional settings, for which the template doesn't provide GUI options.

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 Overview" on page 9 .

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

IN THIS SECTION

NOTE: Ignore the steps in this topic if your switches are already onboarded to the Mist cloud.

To configure and manage a switch through Juniper Mist cloud, you must ensure that the switch is added to the Mist cloud. To see the switch models supported by Mist, visit Juniper Mist Supported Hardware.

You can add greenfield or brownfield switches to the Mist cloud.

In this context, greenfield switches are new cloud-ready switches, while brownfield switches are the switches that are being brought into the Juniper Mist cloud architecture from a previous deployment.

Switch Onboarding Prerequisites

Before you onboard a switch:

- Ensure that you have a Juniper Mist 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 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 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.

Video: Onboard One or More Switches Using a Web Browser

Onboard a Brownfield Switch

It is important to back up your existing Junos OS configuration on the switch before activating a brownfield switch because when the switch is adopted for management from the Juniper Mist cloud, the old configuration is replaced. Back up your existing Junos OS configuration by running the `request system software configuration-backup (path)` command, which saves the currently active configuration and any installation-specific parameters.

In this procedure, you will make a few configuration changes to the Juniper Mist portal, and some to the switch using the Junos OS CLI. Be sure you can log in to both systems.

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 a SSH connection to the Juniper Mist cloud over TCP port 2200 (the switch connection is from a management interface and is used for configuration settings and sending telemetry data).

Figure 1: The Switch Adoption Page

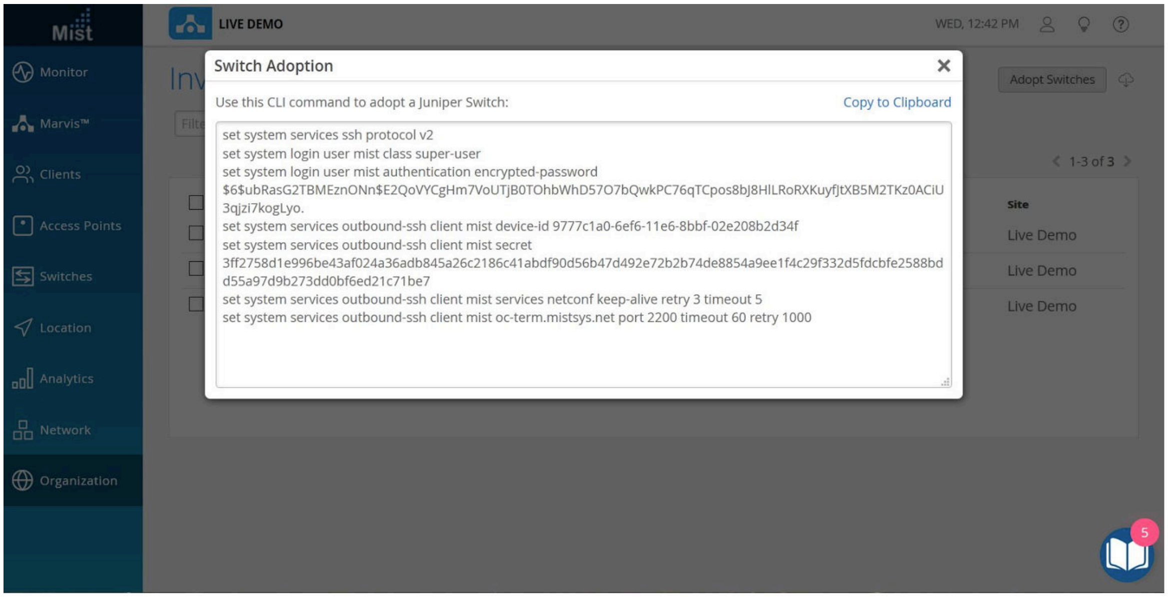

Switch Adoption ✕

Use this CLI command to adopt a Juniper Switch: Copy to Clipboard

```
set system services ssh protocol v2
set system login user mist class super-user
set system login user mist authentication encrypted-password
$6$ubRasG2TBMEznONn$E2QoVYCgHm7VoUTjB0TOhbWhD57O7bQwkPC76qTCpos8bJ8HlLRoRXKuyfJtXB5M2TKz0ACiU
3qjzi7kogLyo.
set system services outbound-ssh client mist device-id 9777c1a0-6ef6-11e6-8bbf-02e208b2d34f
set system services outbound-ssh client mist secret
3ff2758d1e996be43af024a36adb845a26c2186c41abdf90d56b47d492e72b2b74de8854a9ee1f4c29f332d5fdcbfe2588bd
d55a97d9b273dd0bf6ed21c71be7
set system services outbound-ssh client mist services netconf keep-alive retry 3 timeout 5
set system services outbound-ssh client mist oc-term.mistsys.net port 2200 timeout 60 retry 1000
```

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. If you want to add a system message, use the following command:

```
user@host# 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 juniper-mist]` level of the hierarchy.

```
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;
    }
}
```

```
show system login user juniper-mist
```

```
user@Switch-1#  show system login user juniper-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, and then select **Manage configuration with Mist**.

Video: Onboard a Brownfield Switch

Configure Switches

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 process of configuring a switch 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 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.

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: *Template Name* page appears.

   > **NOTE**: You have the flexibility to 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 62 .

3. In the **All Switches Configuration** section, configure basic settings for the switches. Use the tips on the screen to configure the settings.

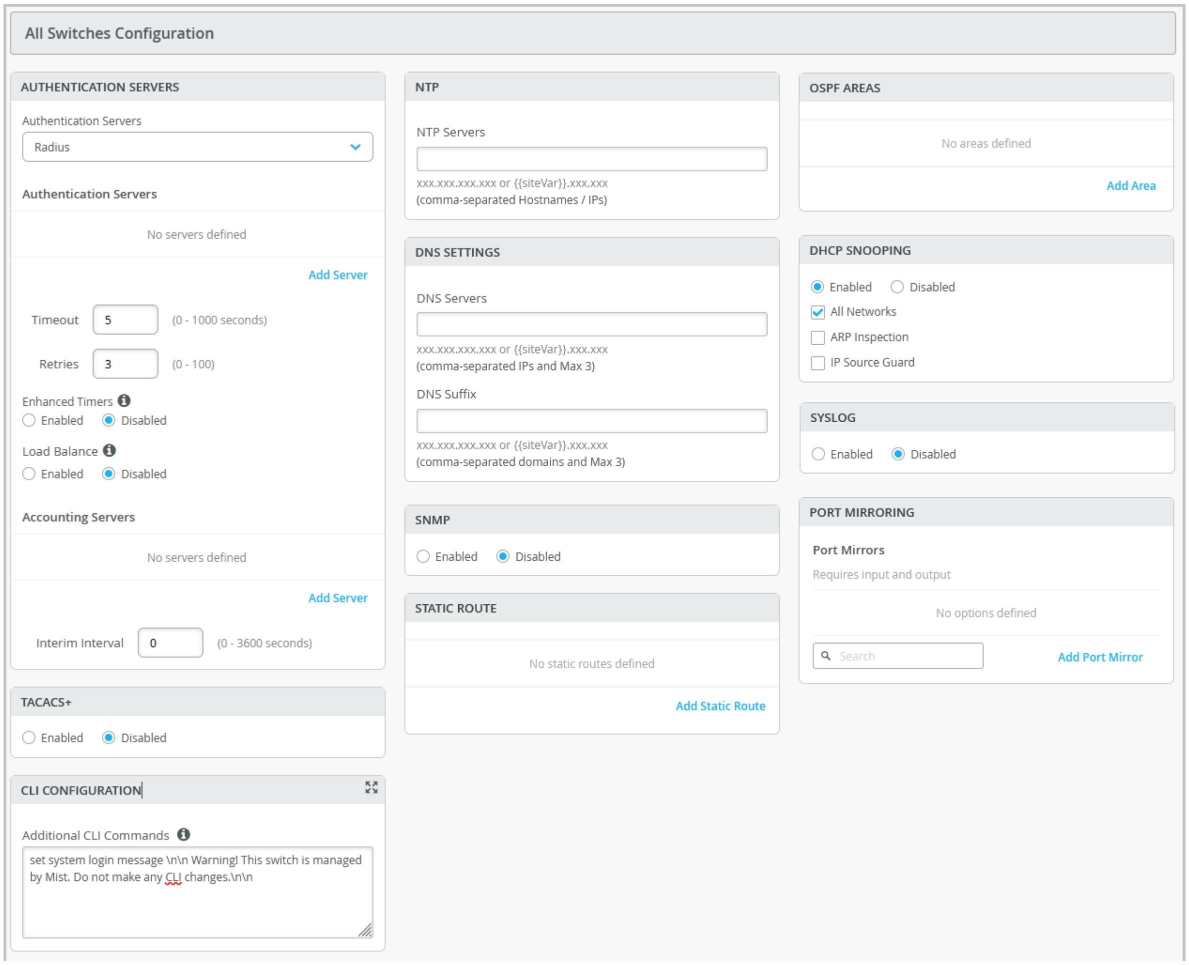

All Switches Configuration

AUTHENTICATION SERVERS

Authentication Servers

Radius

Authentication Servers

No servers defined

Add Server

Timeout 5 (0 - 1000 seconds)

Retries 3 (0 - 100)

Enhanced Timers

○ Enabled ● Disabled

Load Balance

○ Enabled ● Disabled

Accounting Servers

No servers defined

Add Server

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

Add Static Route

OSPF AREAS

No areas defined

Add Area

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

Search

Add Port Mirror

Table 2: All Switches Configuration Field Descriptions

Field	Description
RADIUS	Choose an authentication server for validating usernames and passwords, certificates, or other authentication factors provided by users. • **Mist Auth**—Select this option if you want to configure Juniper Mist Access Assurance, a cloud-based authentication service from Mist, on your switch. For this option to work, you also need to use a port with dot1x or MAB authentication. For more information, see the 'Introducing Mist Access Assurance' section on this product updates page. **NOTE**: Mist Auth on wired switches requires Junos 20.4R3-S7 or above, 22.3R3 or above, 22.4R2 or above, or 23.1R1 or above. To configure Mist Access Assurance features such as authentication policies, policy label, certificates, and identity providers, navigate to **Organization** > **Access**. • **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 **NOTE**: If you want to set up dot1x 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:

Table 2: All Switches Configuration Field Descriptions *(Continued)*

Field	Description
	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*
TACACS+	Configure TACACS+ for centralized user authentication on network devices. Additionally, you can enable TACACS+ accounting on the device to gather statistical data about user logins and logouts on a LAN, and send this data to a TACACS+ accounting server. The port range supported for TACACS+ and accounting servers is 1 to 65535.
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.
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 Field Descriptions *(Continued)*

Field	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) ◦ **Client**—Define a list of SNMP clients. This configuration includes a name for the client list and IP addresses of the clients (in comma separated format). ◦ **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. ◦ **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. • Options under SNMPv3 (V3) ◦ **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. ◦ **VACM**—Define a view-based access control model (VACM). A VACM lets you set access

Table 2: All Switches Configuration Field Descriptions ***(Continued)***

Field	Description
	privileges for a group. You can control access by filtering the MIB objects available for read, write, and notify operations using a predefined view. Each view can be associated with a specific security model (v1, v2c, or usm) and security level (authenticated, privacy, or none). • **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. • Option under both the versions (V2 and V3) • **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. For SNMPv3, you can configure an engine ID, which serves as a unique identifier for SNMPv3 entities. • **View**—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

Table 2: All Switches Configuration Field Descriptions *(Continued)*

Field	Description
	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.
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. Types of static routes supported: • **Subnet**—Includes the IP addresses for the destination network and the next hop. • **Network**—Includes a VLAN (containing a VLAN ID and a subnet) and the next hop IP address.

Table 2: All Switches Configuration Field Descriptions *(Continued)*

Field	Description

CLI CONFIGURATION	For any additional settings that are not available in the template's GUI, you can still configure them using **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 AREAS	Define an Open Shortest Path First (OSPF) area, if required. 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.

Table 2: All Switches Configuration Field Descriptions *(Continued)*

Field	Description
DHCP SNOOPING	Enable the DHCP snooping option to monitor DHCP messages from untrusted devices connected to the switch. DHCP snooping creates a database to keep track of these messages. This helps prevent the acceptance of DHCPOFFER packets on untrusted ports, assuming they originate from unauthorized DHCP servers. DHCP configuration has the following options: • All Networks— Select the All Networks check box to enable DHCP snooping on all VLANs. • Networks—If you want to enable DHCP snooping only on specific networks, click Add (+) in the Networks box and add the required VLANs. • Address Resolution Protocol (ARP) Inspection—Enable this feature to block any man-in-the-middle attacks. ARP Inspection examines the source MAC address in ARP packets received on untrusted ports. It validates the address against the DHCP snooping database. If the source MAC address does not have a matching entry (IP-MAC binding) in the database, it drops the packets. You can check ARP statistics by using the following CLI commands: `show dhcp-security arp inspection statistics`, and `show log messages | match DAI`. The device logs the number of invalid ARP packets that it receives on each interface, along with the sender’s IP and MAC addresses. You can use these log messages to discover ARP spoofing on the network. • IP Source Guard—IP source guard validates the source IP and MAC addresses received on untrusted ports against entries in the DHCP snooping database. If the source addresses do

Table 2: All Switches Configuration Field Descriptions *(Continued)*

Field	Description
	not have matching entries in the database, IP Source Guard discards the packet. **NOTE**: IP Source Guard works only with single-supplicant 802.1X user authentication mode. **NOTE**: • If you have a DHCP server connected to an untrusted access port, DHCP won't function properly. In such cases, you may need to make adjustments to ensure that DHCP works as intended. By default, DHCP considers all trunk ports as trusted and all access ports as untrusted. • You need to enable VLAN on the switch for the DHCP snooping configuration to take effect. So you need to apply port profiles (described later in this document) to the ports. A device with a static IP address might not have a matching MAC-IP binding in the DHCP snooping database, if you have connected the device to an untrusted port on the switch. To check the DHCP snooping database on your switch and view the bindings, use the CLI command `show dhcp-security binding`. This command will provide you with information about the DHCP bindings recorded in the snooping database. For more information, see DHCP Snooping and Port Security Considerations. **NOTE**: You need to enable this feature if you want to view the **DHCP** issues for the switch under the **Successful Connect** SLE metric.

Table 2: All Switches Configuration Field Descriptions *(Continued)*

Field	Description
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. Here are the configuration options available for SYSLOG settings: • **Files**—Send log messages to a named file. • **Hosts**—Send log messages to a remote location. This could be an IP address or hostname of a device that will be notified whenever those log messages are generated. • **Users**—Notify a specific user of the log event. • **Console**—Send log messages of a specified class and severity to the console. Log messages include priority information, which provides details about the facility and severity levels of the log messages. • **Archive**—Define parameters for archiving log messages. • **General**—Specify general information such as a time format, routing instance, and source address for the log messages.

Table 2: All Switches Configuration Field Descriptions *(Continued)*

Field	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. 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 interface to which you want to mirror the traffic. You cannot specify the same interface or network in both the input and output fields.

4. In the **Management** section of the Switch Template Configuration page, configure the following:

   - **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 here.

   - **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 SSH and BGP protocol packets from specified trusted sources. For more information, refer to Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources.

5. In the **Shared Elements** section, configure the following:

   a. In the **Networks** tile, click **Add Network** and configure the VLANs to be used in the port profiles. The settings include a Name, VLAN ID, and a Subnet.

   b. In the **Port Profiles** tile, choose a predefined port profile or click **Add Profile** to create a new profile and assign a network to it. Port profiles provide a way to automate provisioning of multiple switch interfaces. Use the tips on the screen to configure the port profile settings. To know more about port profiles, see "Port Profiles Overview" on page 9 .

Table 3: Key Fields in Port Profile

Field	Description
Name	Name of the port profile.
Mode	Select a port mode. • Trunk—In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches, APs, and routers on the LAN. • Access—Default mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices, such as PCs, printers, IP phones, and IP cameras.

Table 3: Key Fields in Port Profile *(Continued)*

Field	Description
Use dot1x authentication	If you want to use dot1x authentication, select this option. When you select this option, the following additional options are displayed for selection: • MAC authentication—Select this option to set MAC authentication for the port. You can also specify the authentication protocols. Junos supports the following protocols: Protected Extensible Authentication Protocol (eap-peap), Password Authentication Protocol (pap), and eap-md5. These protocols are used by a supplicant to provide authentication credentials for MAC RADIUS authentication. • Allow Multiple Supplicant—Select this option to configure a switch port with dot1x authentication in multiple supplicant mode. In this mode, multiple end-devices can connect to the port. And each end-device is authenticated individually. • Use Guest Network—Select this option to set a guest network to be used for providing authentication. • Bypass authentication when server is down You need to also do the following for dot1x authentication to work: • 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.

Table 3: Key Fields in Port Profile *(Continued)*

Field	Description
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 default value: 0 Supported range: 0 through 16383
PoE	Enable the port to support power over Ethernet (PoE).
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. This setting ensures that the port is treated as an edge port and guards against the reception of BPDUs, which are control messages in the STP. 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 77 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 can also configure STP Edge at the switch level, from the Port Profile section on the switch details page. For more information on STPs, see How Spanning Tree Protocols Work.

Table 3: Key Fields in Port Profile *(Continued)*

Field	Description
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 52 .
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. For more information, see Understanding Storm Control.

Table 3: Key Fields in Port Profile *(Continued)*

Field	Description
Persistent (Sticky) MAC Learning	Enable Persistent (Sticky) MAC to stop unauthorized devices from connecting to your network. When enabled, the switch learns the MAC addresses of devices that arrive on the port and saves them in memory. If the number of MAC addresses learned exceeds the 'MAC Limit' specified above, the port drops the frames. Also, you will see a 'MAC Limit Exceeded' event on the Insights page. You can hover over the port from the front panel on the switch details page to see the MAC Limit and the MAC Count (the number of MAC addresses that the port learned dynamically). **NOTE**: • You cannot enable this feature on a Trunk port or on a port with 802.1X authentication, as Junos OS does not support this combination. • Enable this feature for static wired clients. Do not enable it for Mist AP interfaces. The Juniper Mist portal does not show the MAC addresses that an interface has learned. It shows only the maximum MAC address count. To view the MAC addresses that an interface learned, select the **Utilities** > **Remote Shell** option on the switch details page and run the following commands: • `show ethernet-switching table persistent-learning` • `show ethernet-switching table persistent-learning interface` The MAC Count value remains on the port until you clear it from the front panel on the switch details or until you disable the Persistent (Sticky) MAC Learning feature. To clear the MAC addresses that a port learned, select the port on the switch front panel and then click **Clear MAC [Dynamic/Persistent]**. This action generates a MAC Limit

Table 3: Key Fields in Port Profile *(Continued)*

Field	Description
	Reset event on the Switch Insights page. Read more about the front panel in "Switch Details" on page 77 .

c. In the **VRF** tile, configure Virtual Routing and Forwarding (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.

   NOTE: You can't assign the default network (VLAN ID = 1) to VRF.

d. In the **Dynamic Port Configuration** tile, set up rules for dynamically assigning port profiles. When a user connects a client device to a switch port with this feature enabled, 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 pre-configured port and network settings to the interface. You can configure a dynamic port profile based on the following parameters:

   - LLDP System Name
   - LLDP Description
   - LLDP Chassis ID
   - Radius Username
   - Radius Filter-ID
   - MAC (Ethernet mac-address)

   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, it assigns the 'AP' profile to the connected device.

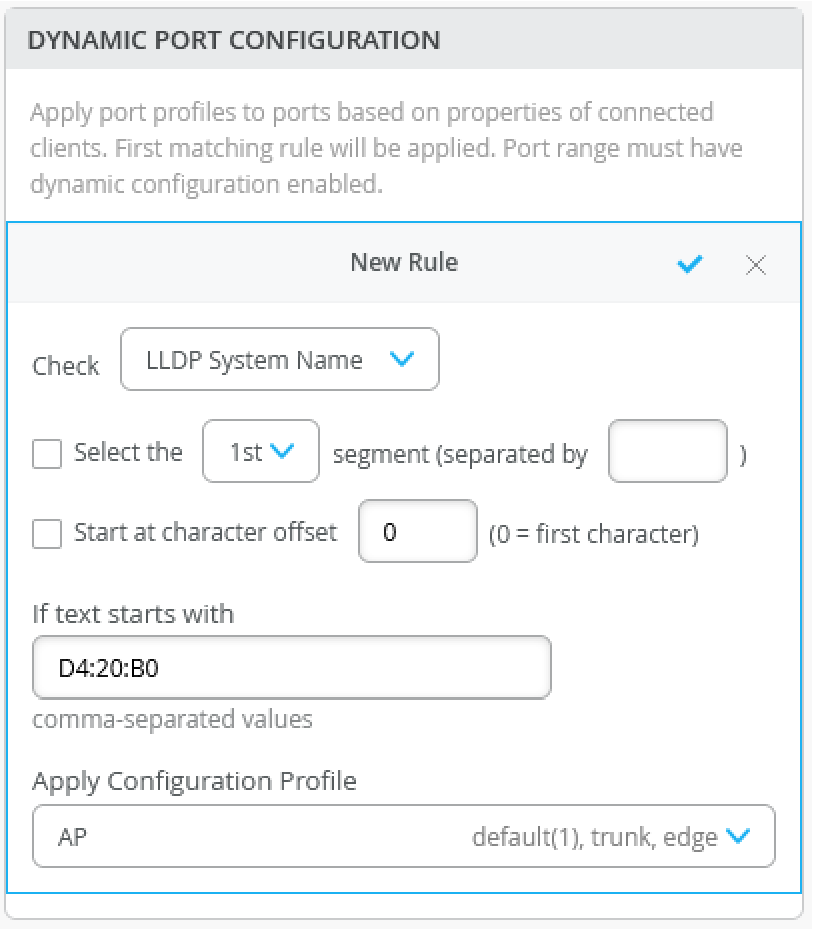

DYNAMIC PORT CONFIGURATION

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

comma-separated values

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

For your dynamic port configurations to take effect, you also need to 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. You can also do this at the switch level, from the Port Configuration section on the switch details page.

NOTE:

- It takes a couple 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).

6. In the **Select Switches Configuration** section, configure the following:

   a. On the **Info** tab, create a rule to associate the shared elements with your switch. Here's an example of how to add a rule that maps the EX4300 switch to an "access" role.

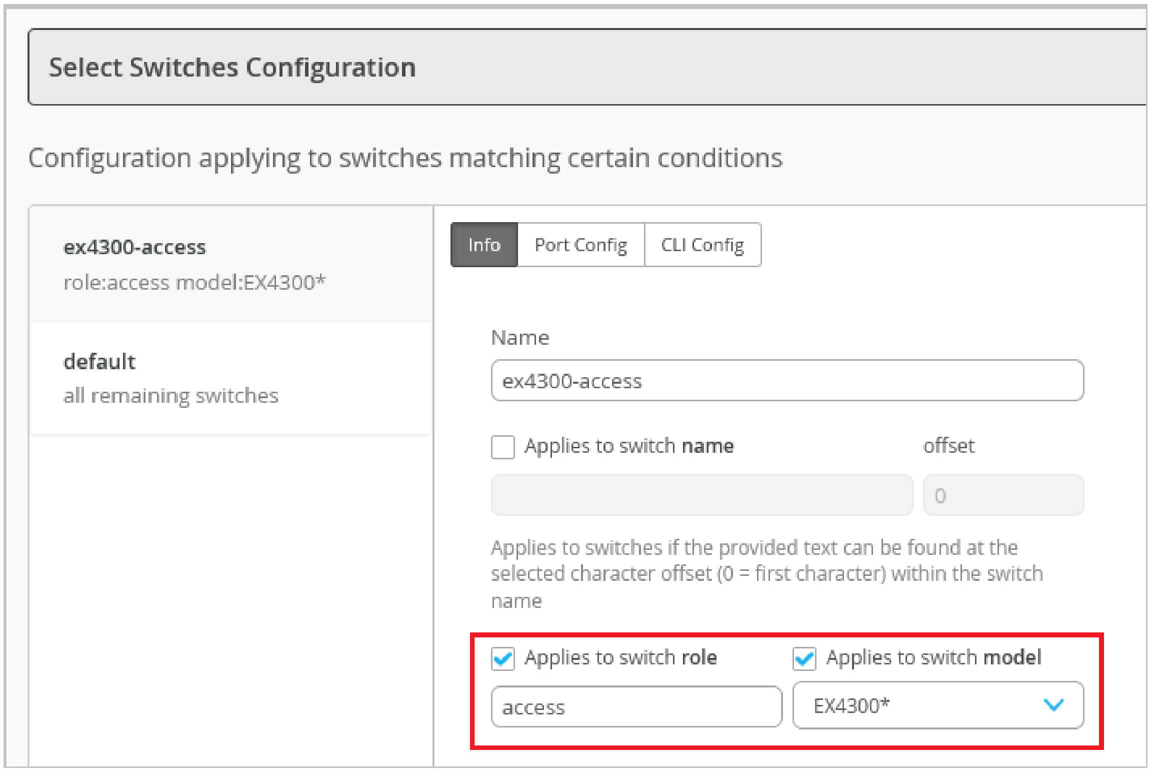

Select Switches Configuration

Configuration applying to switches matching certain conditions

ex4300-access
role:access model:EX4300*

default
all remaining switches

Name: ex4300-access

☐ Applies to switch **name** 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**: access

☑ Applies to switch **model**: EX4300*

b. On the Port Config tab, click **Add Port Range** to associate a port profile with a port. Here you also have the following key options:

- Enable Dynamic Configuration on the port. Dynamic port profiling allows you to assign a dynamic profile to a connected device based on defined attributes. 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 will be placed in a specified VLAN. 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.

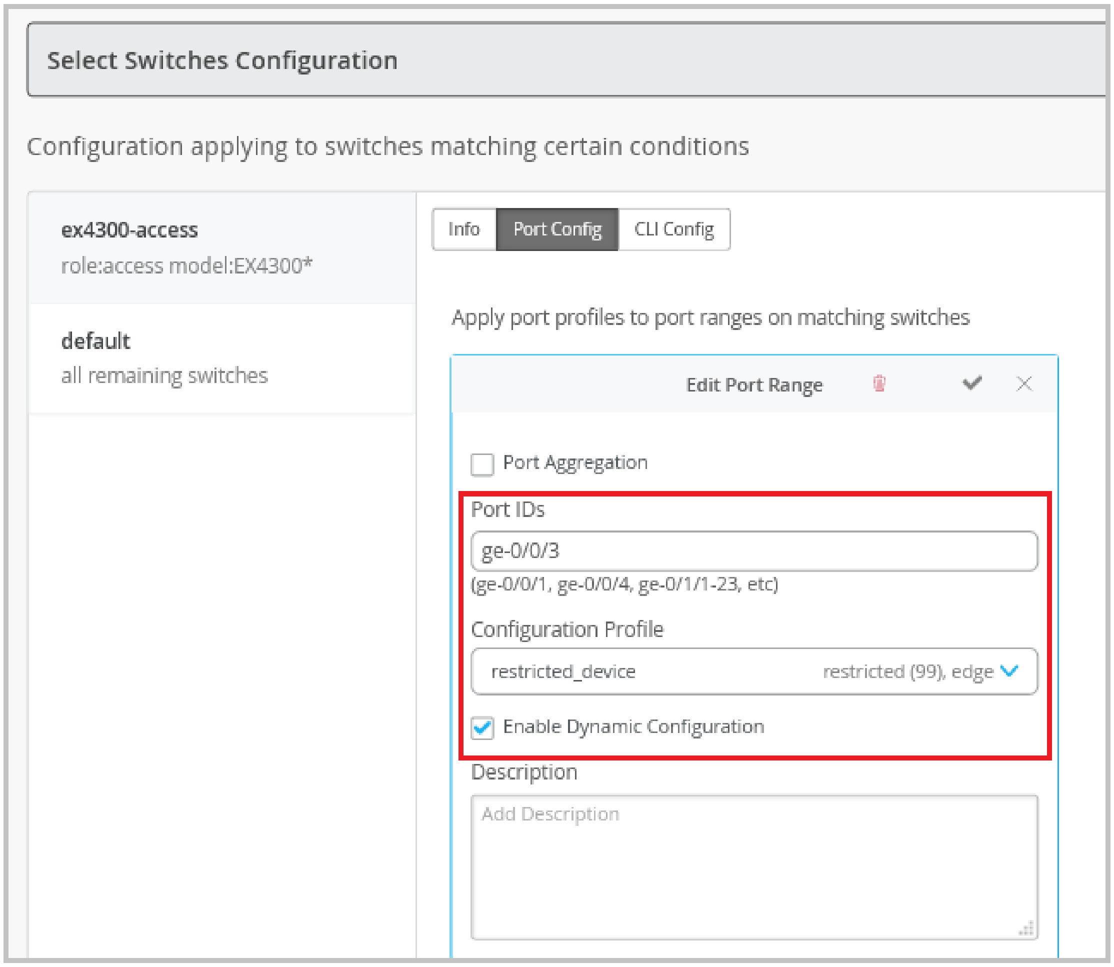

Select Switches Configuration

Configuration applying to switches matching certain conditions

ex4300-access
role:access model:EX4300*

default
all remaining switches

Apply port profiles to port ranges on matching switches

Edit Port Range

□ Port Aggregation

Port IDs: ge-0/0/3
(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Configuration Profile: restricted_device — restricted (99), edge

☑ Enable Dynamic Configuration

Description: Add Description

- Enable Port Aggregation. Port aggregation or link aggregation enables you to group Ethernet interfaces 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 configure 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. You can also configure 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. The following example shows the use of LAG in an uplink port configuration:

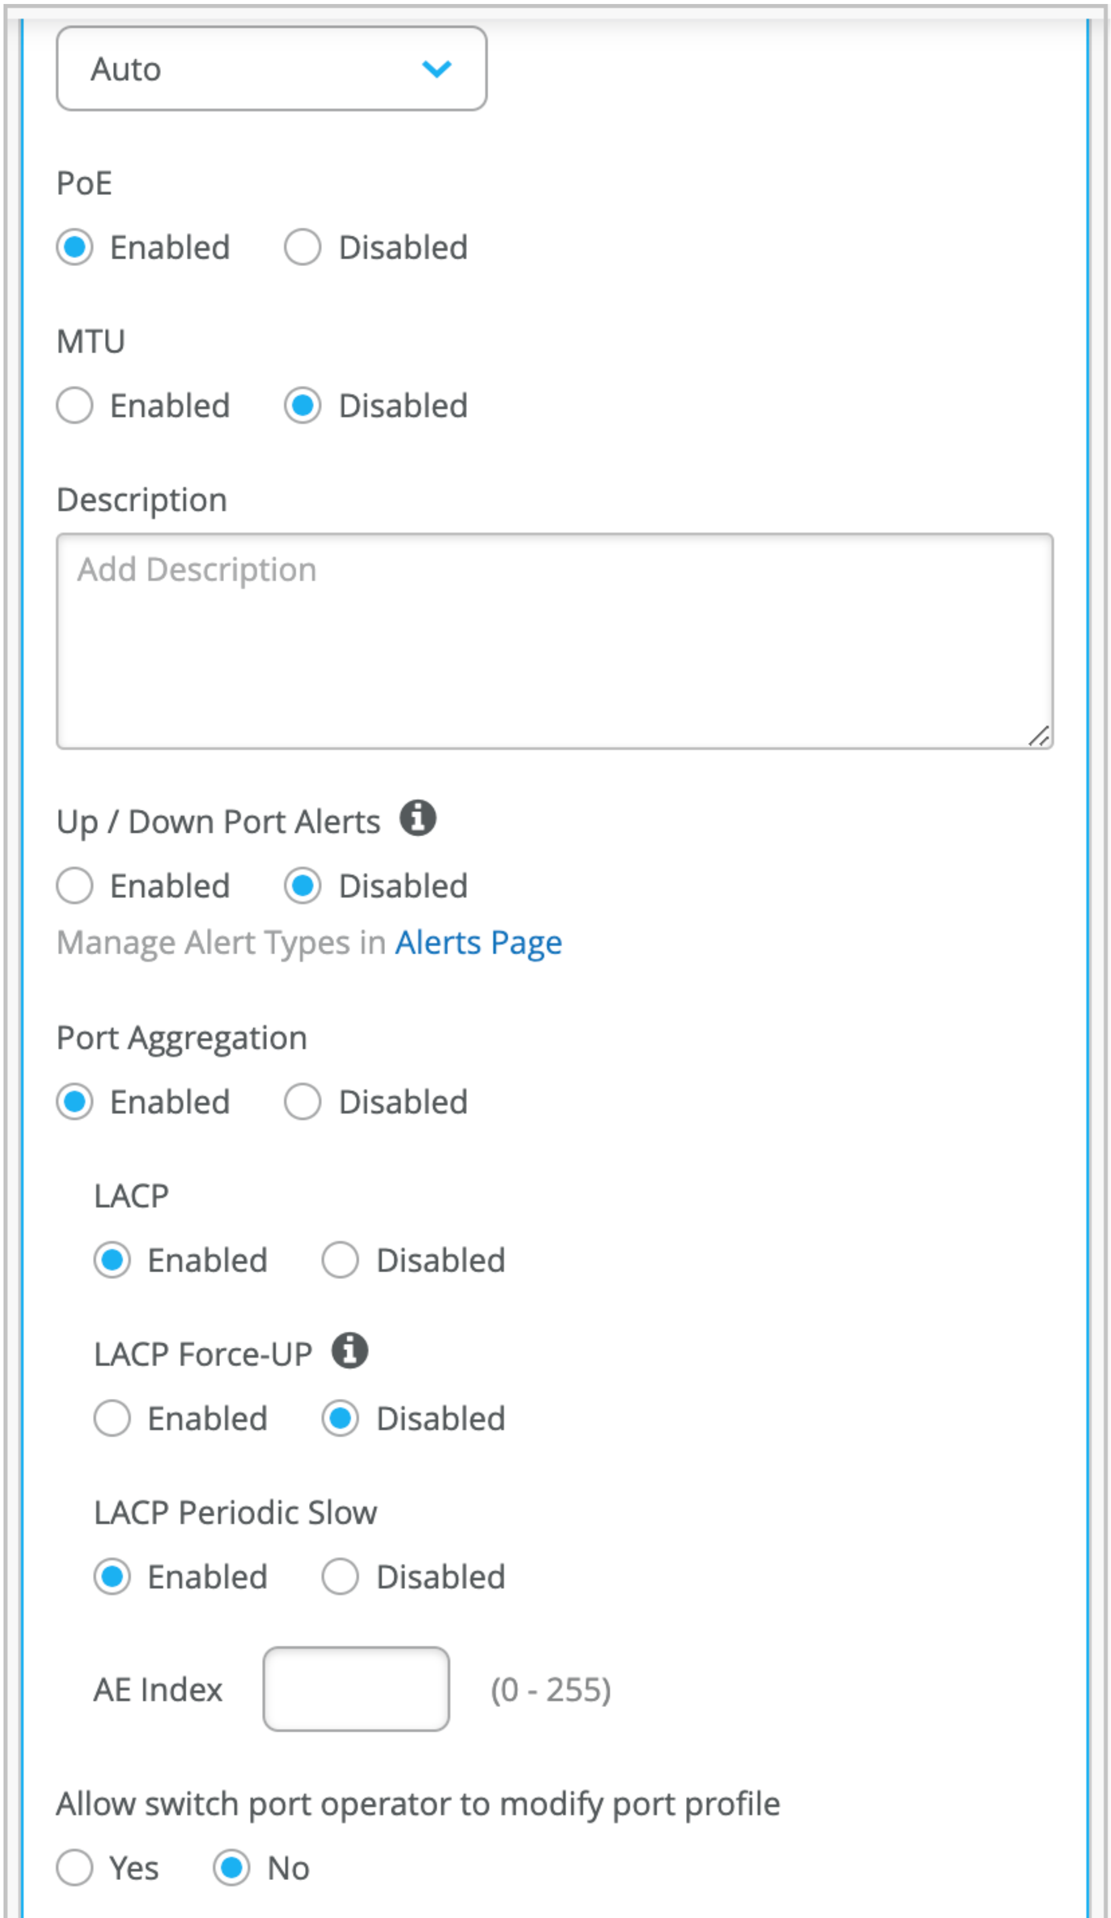

Auto

PoE

● Enabled ○ Disabled

MTU

○ Enabled ● Disabled

Description

Add Description

Up / Down Port Alerts

○ Enabled ● Disabled

Manage Alert Types in Alerts Page

Port Aggregation

● Enabled ○ Disabled

LACP

● Enabled ○ Disabled

LACP Force-UP

○ Enabled ● Disabled

LACP Periodic Slow

● Enabled ○ Disabled

AE Index (0 - 255)

Allow switch port operator to modify port profile

○ Yes ● No

- Configure alerts and email notifications for the interface up and down events on specified ports of a switch. To configure a switch port to support alerts, select the **Enable "Up/Down Port" Alerts** check box. Also, on the **Monitor** > **Alerts** > **Alerts Configuration** page, you must select from the following check boxes to enable alerts for the ports.
  - Critical WAN Edge Port Up
  - Critical WAN Edge Port Down
  - Critical Switch Port Up
  - Critical Switch Port Down

c. On the **CLI Config** tab, include CLIs (in the `set` format) to configure any additional rule-based settings for which the template doesn't provide a GUI option.

7. In the Switch Policy section, configure Group Based Policies (GBPs) that you can use in your campus fabric IP Clos deployments. 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.

   Only the following devices that run Junos OS Release 22.4R1 and later support GBPs: EX4400, EX4100, EX4650, QFX5120-32C and QFX5120-48Y.

   The following image shows a sample GBP:

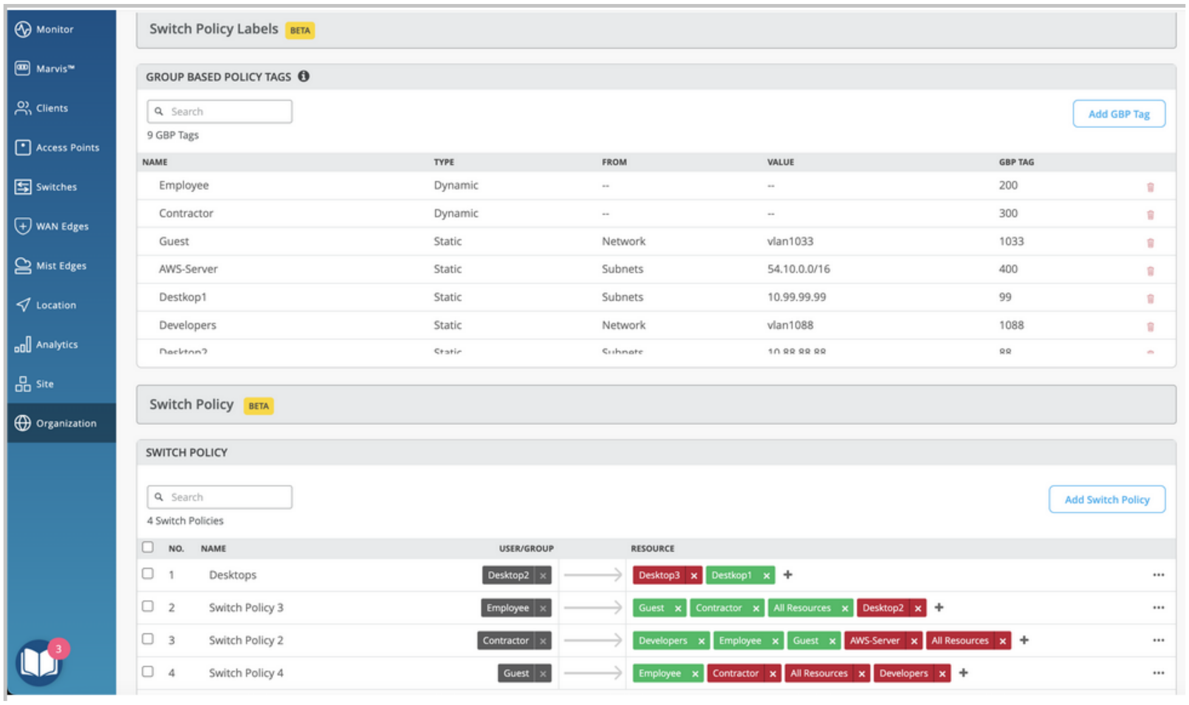

Switch Policy Labels BETA

GROUP BASED POLICY TAGS

Search
Add GBP Tag

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

Switch Policy BETA

SWITCH POLICY

Search
Add Switch Policy

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 × +

To configure GBP:

a. In the **Group Based Policy Tags** section, create a GBP tag. as described below:

i. Click **Add GBP tag**.

ii. Specify a name for the tag.

iii. Choose the tag type—**Dynamic** or **Static**. By default, Juniper Mist chooses the Dynamic option. If you choose the static tag, specify a GBP tag source. It can be a MAC address, network, or an IP subnet.

> **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.

iv. Specify a GBP tag value or GBP source tag for host-originated packets (range: 1 through 65535).

b. In the **Switch Policy** section, add a policy. The policy filters use GBP source tags, destination tags, or both as matches to either allow or discard traffic. To create a policy, use the steps below:

i. Click **Add Switch Policy**.

ii. In the **USER/GROUP** column, click Add (**+**) and add the users or groups that need access to the resources. You can use the GBT tags here, if you have defined them already.

iii. In the **RESOURCE** column, click Add (**+**) and add the resources that you need to map to the selected users or groups. You can use the GBT tags here too, if you have defined them already.

By default, users are given access to the resources added. If you want to deny the user access to certain resources, click the Resource label that you have added and set the access to **deny**. See below:

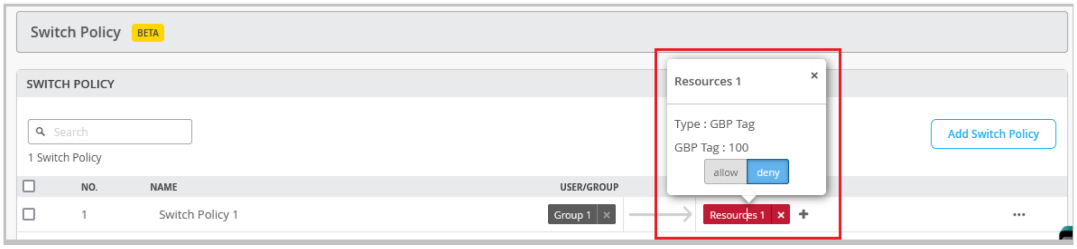

Type : GBP Tag

GBP Tag : 100

8. Click **Save** to save the switch template.

   The **Confirm changes** window appears.

9. 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 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**.
2. From the **List** tab, click the switch you want to edit.
3. Configure the switch-specific settings that include the following:
   - Name—A hostname for the switch.
   - Role—The role of the switch in the network. Example: Access.
   - IP Configuration (Out of Band)—The management interface settings.
   - IP Configuration—The IRB interface settings for inter-VLAN routing.
4. If you want to override the template settings applied to the switch, follow these steps:
   a. Select the **Override Site/Template Settings** 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**.
2. From the **List** tab, 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.

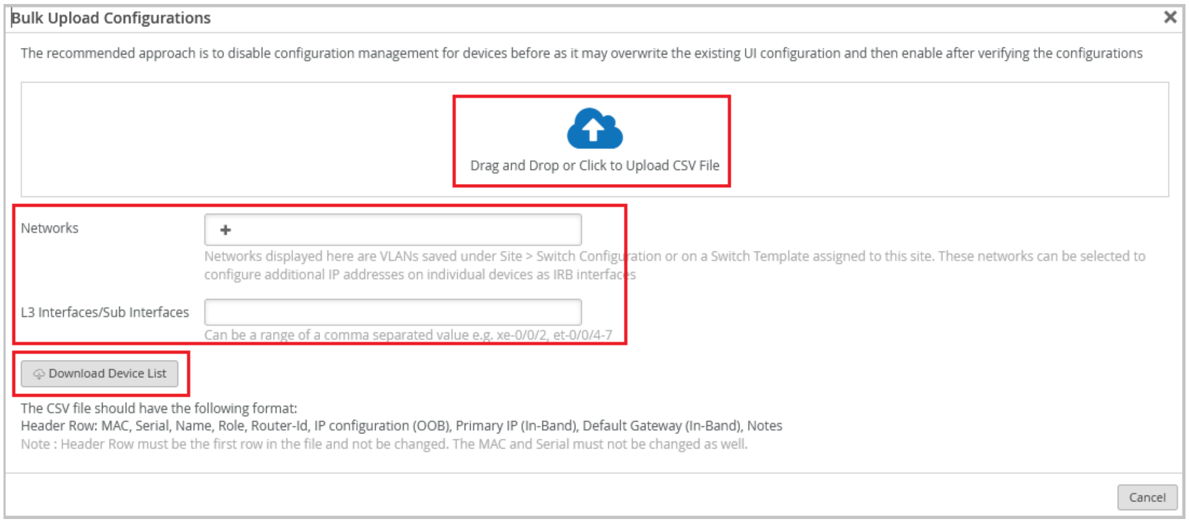

Bulk Upload Configurations

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

Drag and Drop or Click to Upload CSV File

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

Download Device List

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,

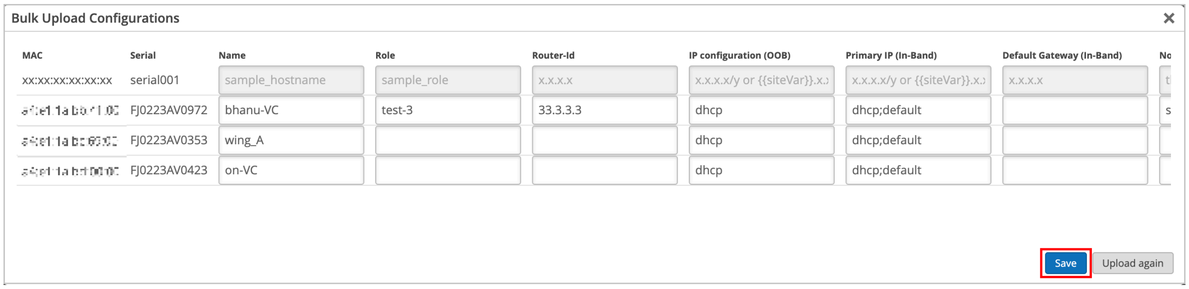

Bulk Upload Configurations

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 the Switch Configuration

You can easily review the configuration applied to your switches and make any updates through the "switch details" on page 77 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.

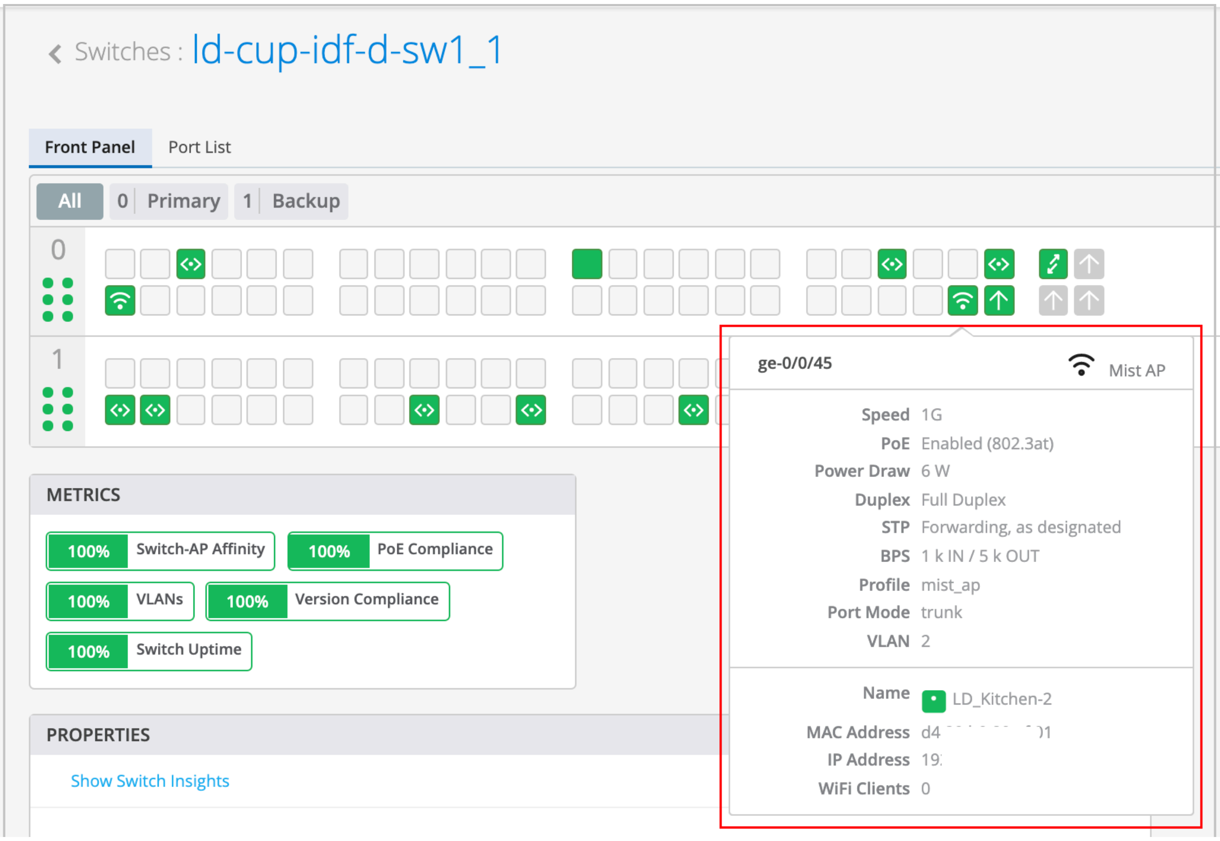

100%	Switch-AP Affinity
100%	PoE Compliance
100%	VLANs
100%	Version Compliance
100%	Switch Uptime

ge-0/0/45 Mist AP

Speed	1G
PoE	Enabled (802.3at)
Power Draw	6 W
Duplex	Full Duplex
STP	Forwarding, as designated
BPS	1 k IN / 5 k OUT
Profile	mist_ap
Port Mode	trunk
VLAN	2
Name	LD_Kitchen-2
MAC Address	d4 [redacted] 01
IP Address	19 [redacted]
WiFi Clients	0

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 82 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 62 .

SEE ALSO

Manage or Update Configuration Settings | 62

Switch Details | 77

Wired Service-Level Expectations (SLEs) | 154

QoS Configuration

IN THIS SECTION

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. The QoS implementation in Juniper Mist™ generally involves the following:

- 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	5	10	10	Strict-high
nc1/Low nc2/Low	net-control	7	3	3	Low

As shown in the preceding 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.

Configure QoS

You can configure QoS on a switch port from the Port Profile tile on the switch details page or switch template.

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.

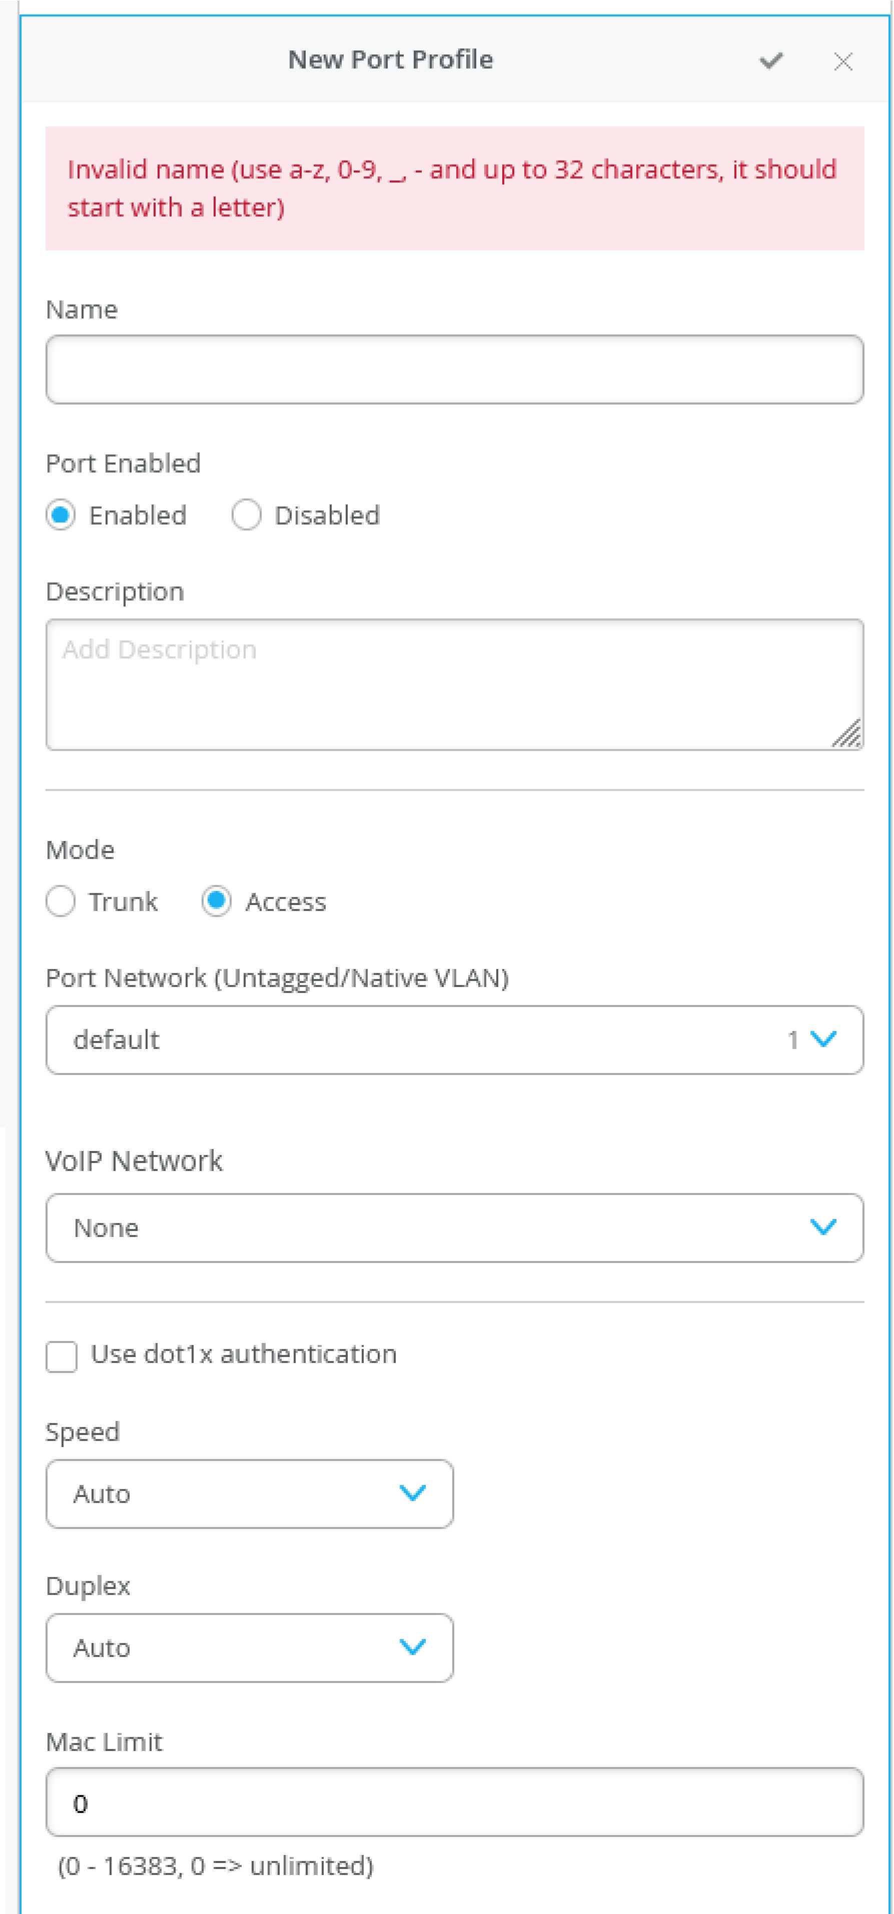

New Port Profile

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)

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.

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

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. The downstream traffic (AP > client) gets marked with the override access class value specified. The override configuration doesn't support upstream traffic (client > AP).

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

Manage or Update Configuration Settings

SUMMARY

You can manage configuration settings at the template level, site level, and device level.

IN THIS SECTION

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.

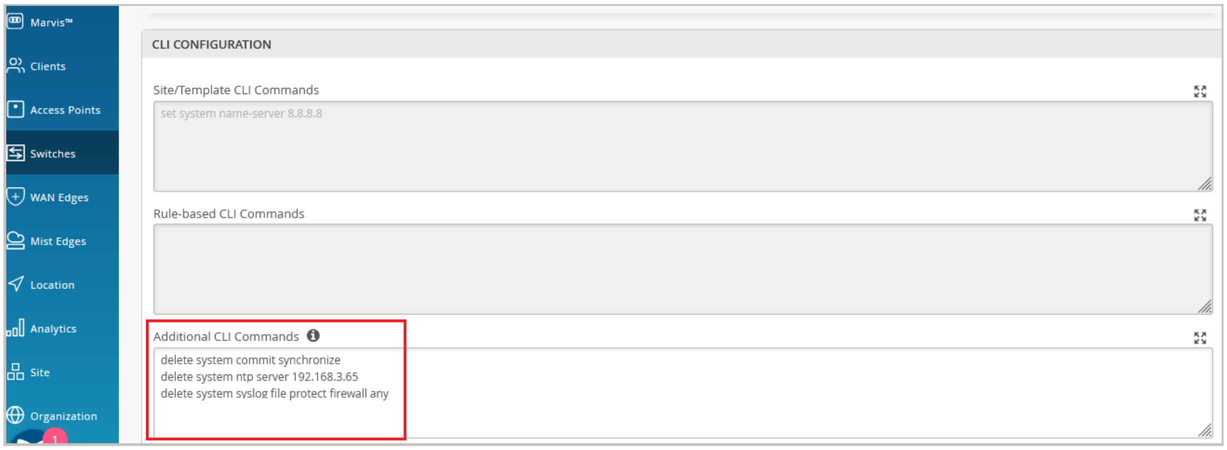

CLI CONFIGURATION

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 77 .
- 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 Software on Your Switch

IN THIS SECTION

You can upgrade the Junos OS version running on your switch from the Juniper Mist™ portal.

Before upgrading the Junos OS software running on your switch, ensure that the switch has the following:

- The storage space required to accommodate the new image.
- A stable SSH connection to the Mist cloud.

Free Up Storage Space on Your Switch

When you initiate the 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 Junos OS Software on Your Switch

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

In the case of Virtual Chassis, you can only upgrade the mixed EX4300 Virtual Chassis, which combines EX4300 multigigabit model (EX4300-48MP) switches with any other EX4300 model switches. Juniper Mist does not support nonstop software upgrade (NSSU).

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 only, 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.

To upgrade the Junos OS software on your switch:

1. Click **Switches** on the left navigation pane in the Juniper Mist portal.
2. 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" on page 161 .
3. From the **List** tab, select the switches that require a software upgrade, and then click **Upgrade Switches**. You can select one or more switches for the upgrade.

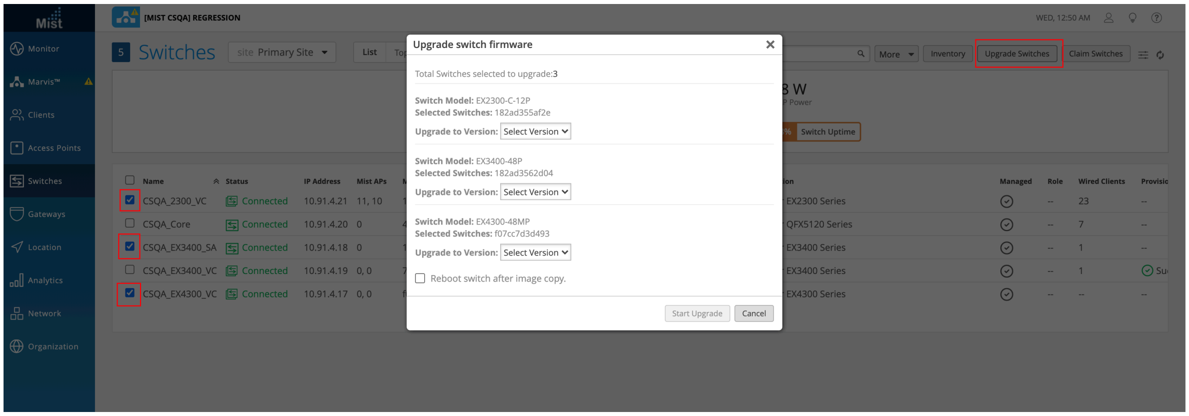

Upgrade switch firmware ×

Total Switches selected to upgrade:3

Switch Model: EX2300-C-12P
Selected Switches: 182ad355af2e
Upgrade to Version: Select Version

Switch Model: EX3400-48P
Selected Switches: 182ad3562d04
Upgrade to Version: Select Version

Switch Model: EX4300-48MP
Selected Switches: f07cc7d3d493
Upgrade to Version: Select Version

☐ Reboot switch after image copy.

Start Upgrade | Cancel

	Name	Status	IP Address	Mist APs	ion	Managed	Role	Wired Clients	Provisio
☑	CSQA_2300_VC	Connected	10.91.4.21	11, 10	EX2300 Series	✓	--	23	--
☐	CSQA_Core	Connected	10.91.4.20	0	QFX5120 Series	✓	--	7	--
☑	CSQA_EX3400_SA	Connected	10.91.4.18	0	EX3400 Series	✓	--	1	--
☐	CSQA_EX3400_VC	Connected	10.91.4.19	0, 0	EX3400 Series	✓	--	1	✓ Su
☑	CSQA_EX4300_VC	Connected	10.91.4.17	0, 0	EX4300 Series	✓	--	--	--

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 77).

4. In the Upgrade Switch Firmware window, select the target software version from the **Upgrade to Version** drop-down list, and then click **Start Upgrade**. 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. Navigate to the switch details page (**Switches** > ***Switch Name***).

     b. Reboot the switch from **Utilities** > **Reboot Switch**.

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.

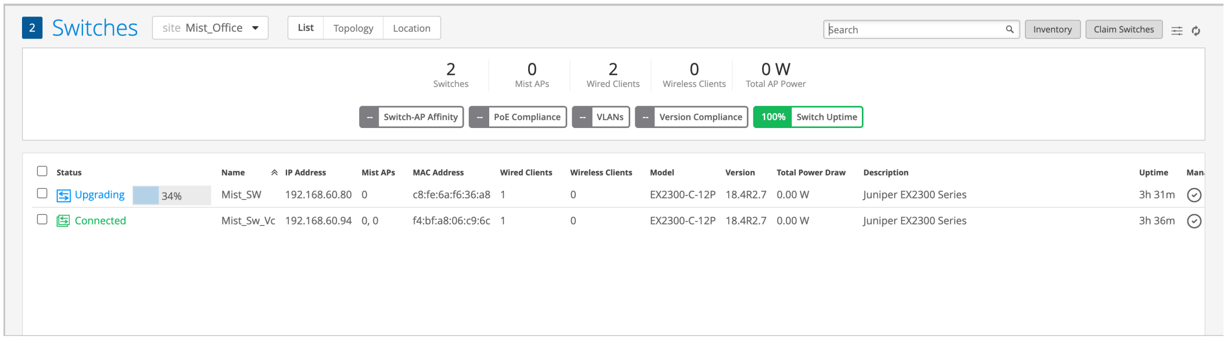

	Status	Name	IP Address	Mist APs	MAC Address	Wired Clients	Wireless Clients	Model	Version	Total Power Draw	Description	Uptime	Man
□	Upgrading 34%	Mist_SW	192.168.60.80	0	c8:fe:6a:f6:36:a8	1	0	EX2300-C-12P	18.4R2.7	0.00 W	Juniper EX2300 Series	3h 31m	✓
□	Connected	Mist_Sw_Vc	192.168.60.94	0, 0	f4:bf:a8:06:c9:6c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

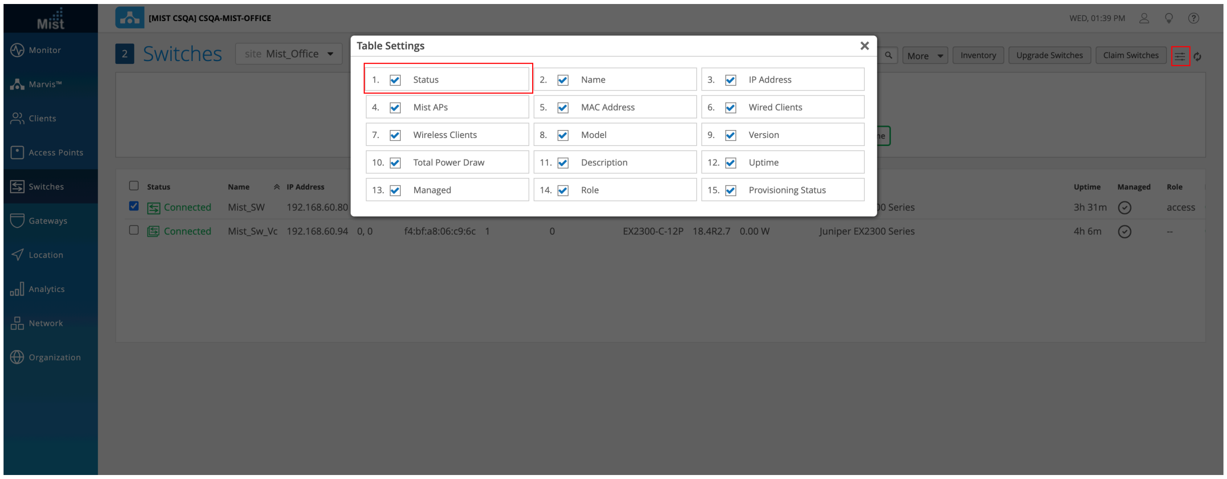

Table Settings

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

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 77 page and click the **Switch Insights** link on the **Properties** tile.

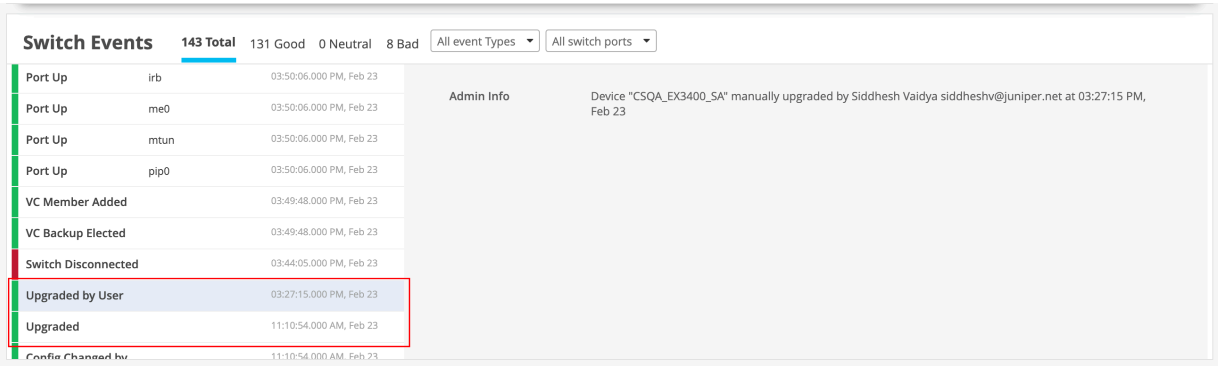

Switch Events 143 Total 131 Good 0 Neutral 8 Bad All event Types All switch ports

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:

  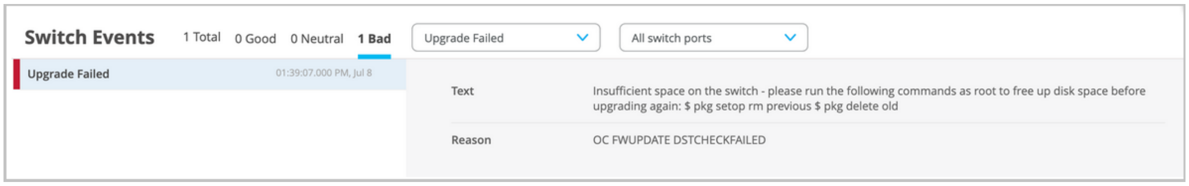

  Switch Events 1 Total 0 Good 0 Neutral 1 Bad Upgrade Failed All switch ports

  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" on page 66

- 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.

Replace a Switch

You can replace a switch in your Juniper Mist™ network from the switch details page 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 77 page appears.
3. Select the **Replace Switch** option from the **Utilities** drop-down list on the details.

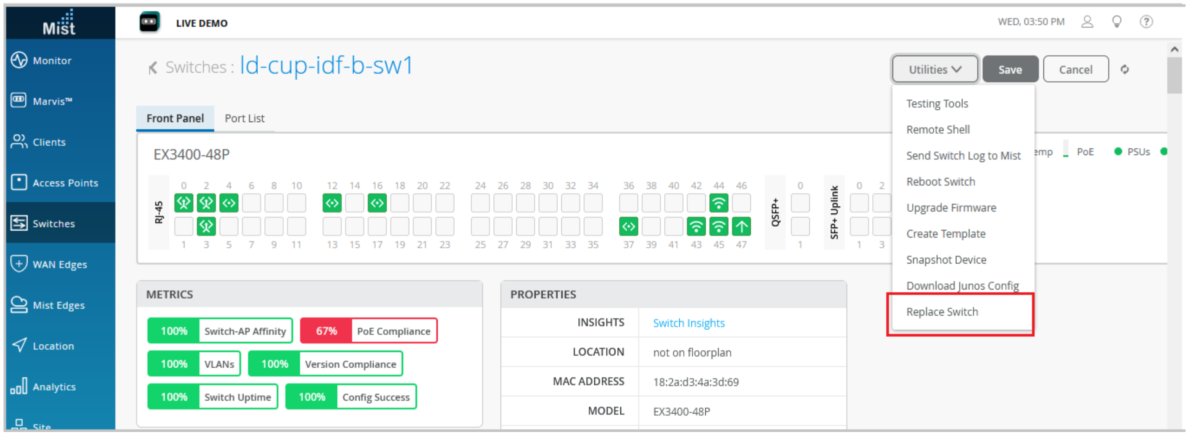

INSIGHTS	Switch Insights
LOCATION	not on floorplan
MAC ADDRESS	18:2a:d3:4a:3d:69
MODEL	EX3400-48P

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.

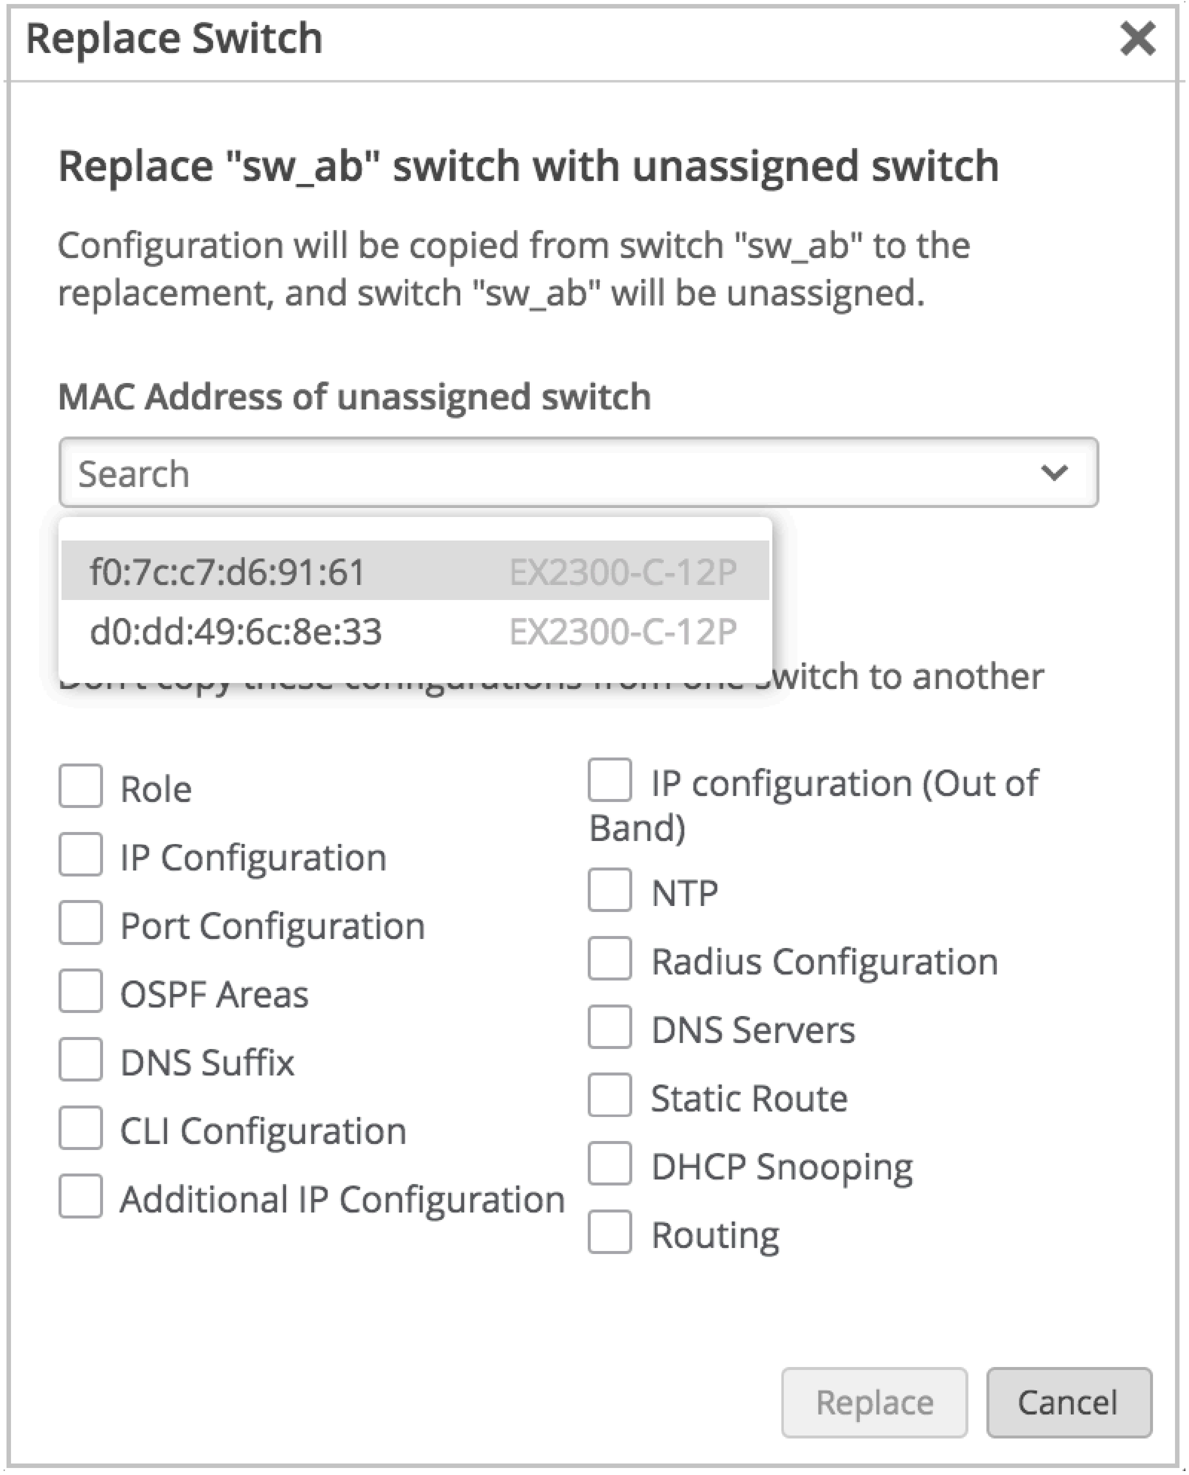

Replace Switch

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" on page 22 .

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.

3

CHAPTER

Switch Dashboards

Switch Metrics

The metrics on the Switches page help you track the switch performance against certain compliance parameters.

Video: Switch Metrics Overview

To view switch metrics, click **Switches** on the left navigation pane on the Mist portal.

Figure 2: Switch Metrics on the Switches Page

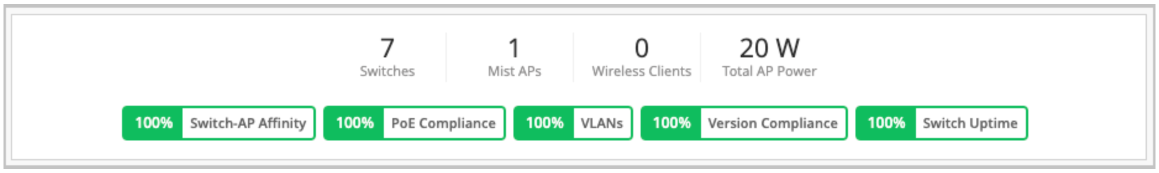

Here's a list of the switch metrics you can track:

- **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.

Switch Details

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.

To access a switch details page:

1. Open the Mist portal and click **Switches** on the left pane.

2. Select a switch from the **List** tab.

   You can use the **Site** drop-down list to filter the switches by a site.

Here's an example of the details page of a switch named **ld-cup-idf-a-sw22**.

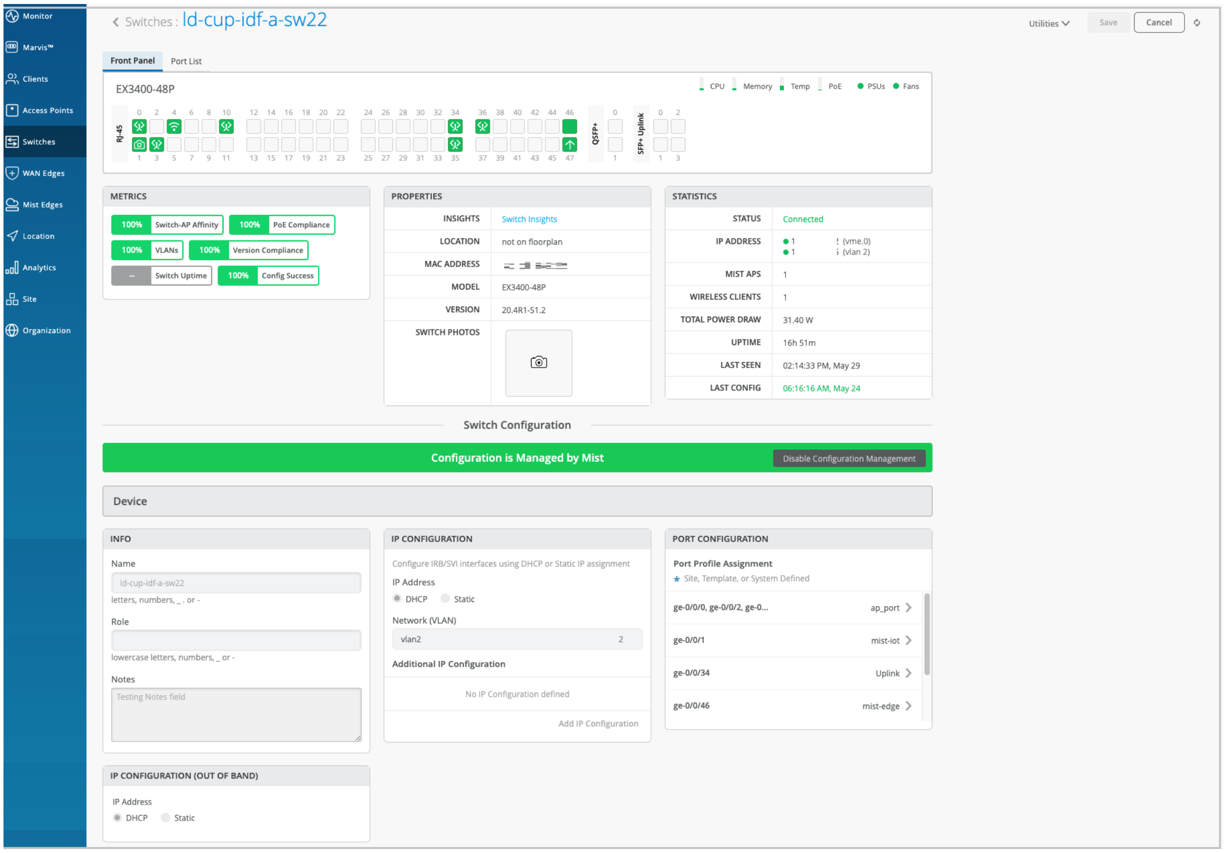

Switches : ld-cup-idf-a-sw22

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]2 (vme.0) ● 1[redacted]5 (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.

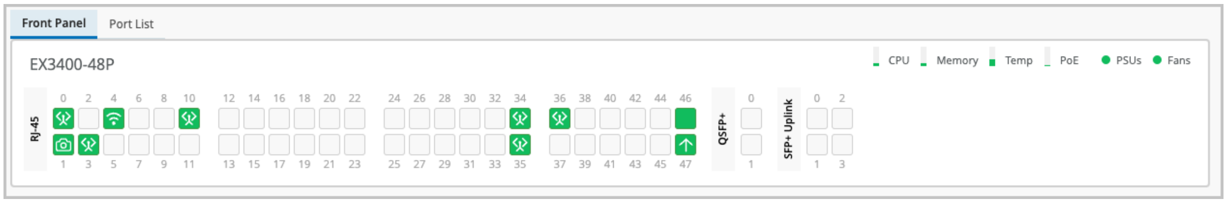

The port icons on the Front Panel view help you quickly identify the client devices or APs connected to each port and their status. The following table lists the key port icons and their descriptions:

Table 4: 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 4: Port Icons and Their Descriptions *(Continued)*

Port Icon	Description
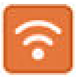	A Mist AP is connected, but the port has active alerts.
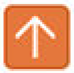	An uplink is connected. But the port has active alerts.

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, and much more.

Here's an example of what you might see if you hover over port ge-0/0/45:

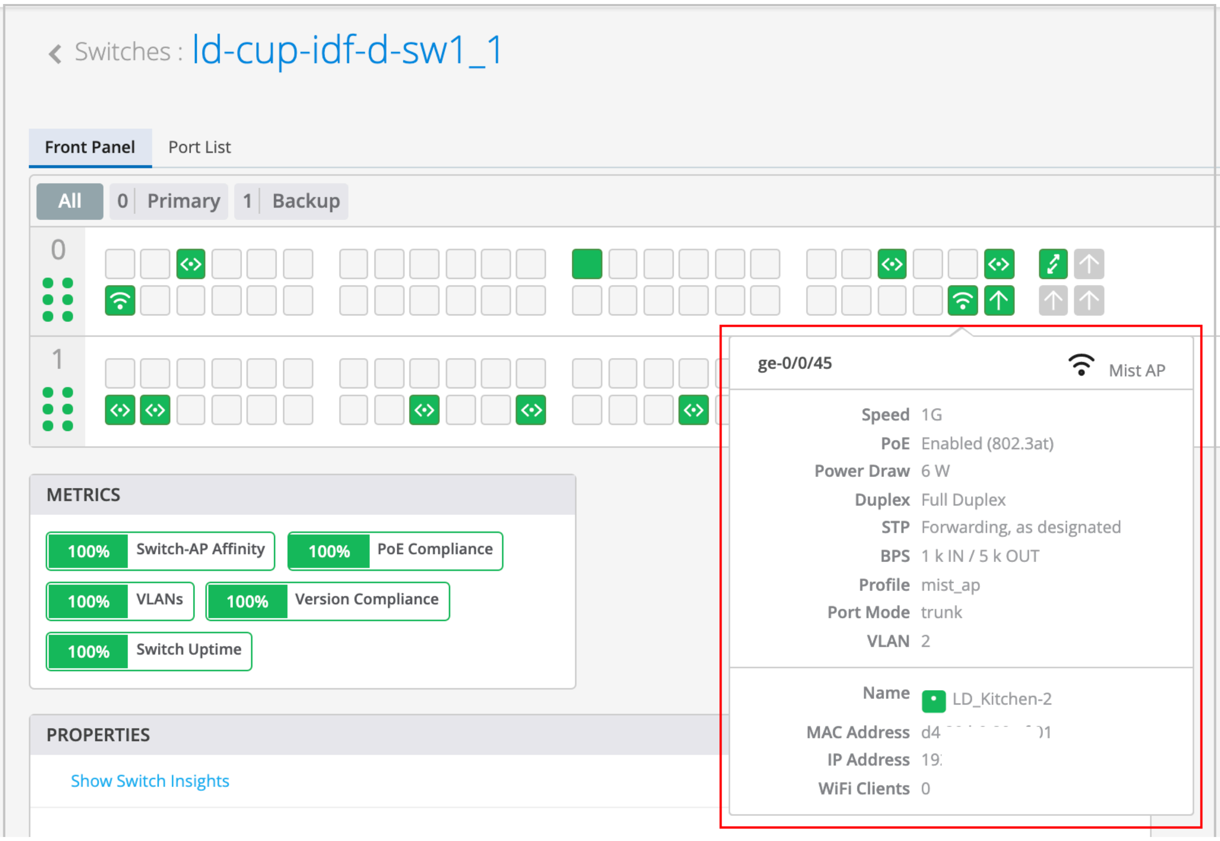

< Switches : ld-cup-idf-d-sw1_1

METRICS

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

PROPERTIES

Show Switch Insights

ge-0/0/45 Mist AP

Speed 1G
PoE Enabled (802.3at)
Power Draw 6 W
Duplex Full Duplex
STP Forwarding, as designated
BPS 1 k IN / 5 k OUT
Profile mist_ap
Port Mode trunk
VLAN 2

Name LD_Kitchen-2
MAC Address d4 [redacted])1
IP Address 19. [redacted]
WiFi Clients 0

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.

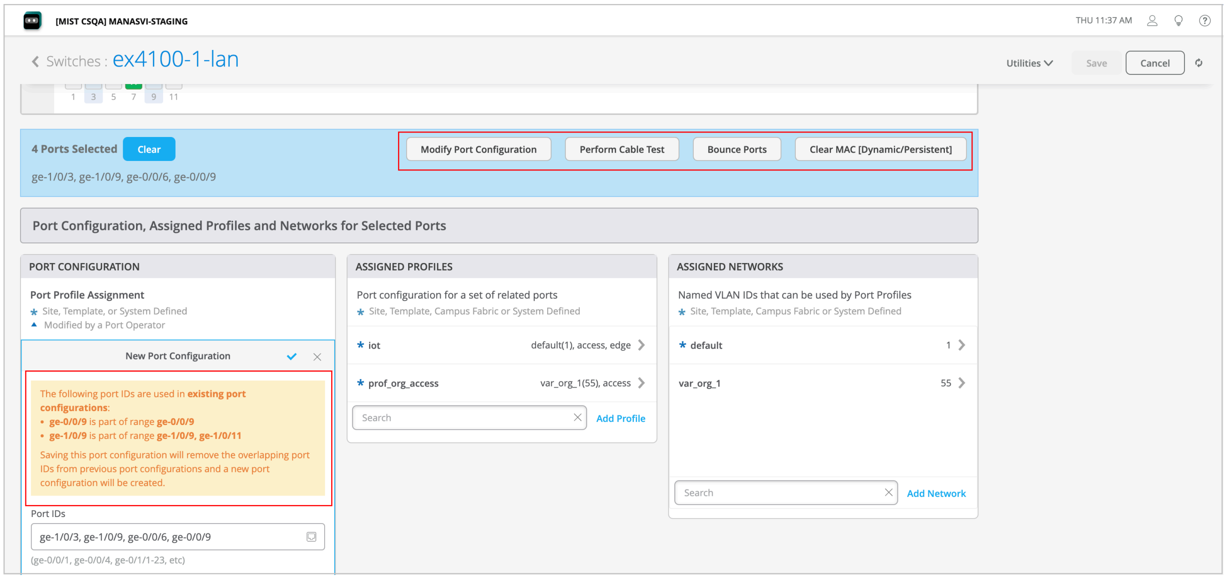

4 Ports Selected

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

New Port Configuration

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.

On the right side of the front panel, you can also see the device usage indicators. To view the usage information, hover over the indicators on the upper right of the screen. For example, you can see the temperature values of each component that contribute to the total temperature of the switch.

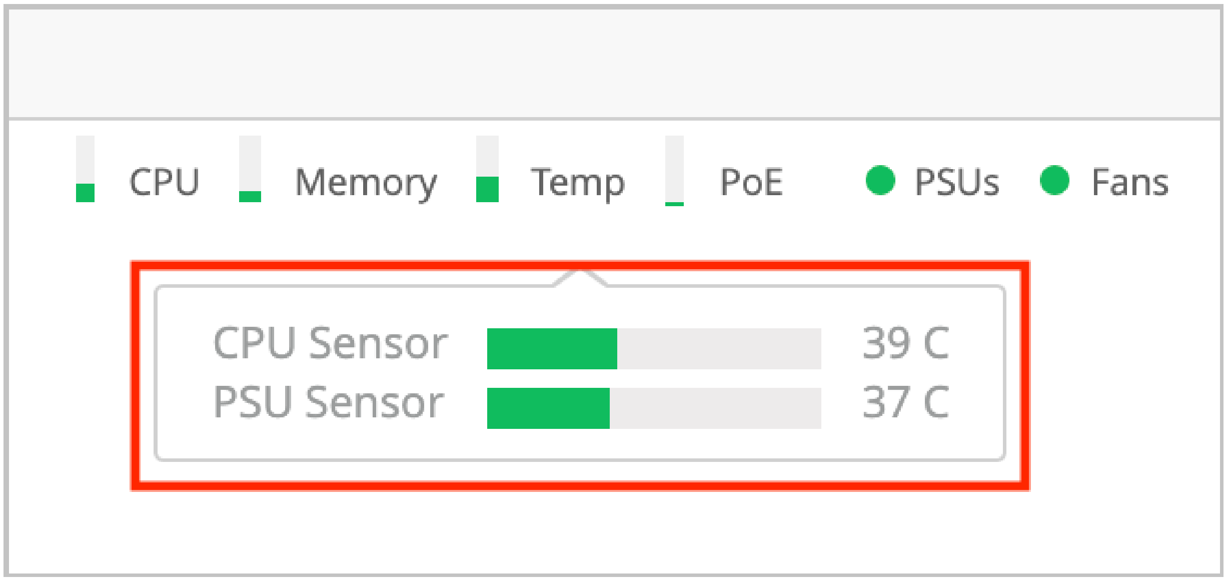

CPU Sensor 39 C
PSU Sensor 37 C

All in all, this detailed view and the additional features make managing your switch a breeze.

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.

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 also can 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.

Switch Insights also gives you a detailed port list. You can see all the ports along with information about the connected endpoints.

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 76 .

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 84 .

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

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**. You can also 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.
- **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.
- **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
```

- **Create Template**—Create a switch configuration template based on the switch configuration.
- **Snapshot Device**—Store a recovery snapshot in the OAM (Operations, Administrations and Maintenance) volume, which holds a full backup that can be used in case something goes wrong with the switch configuration.
- **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.

> **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" on page 22 .

RELATED DOCUMENTATION

Switch Details | 77

4

CHAPTER

Virtual Chassis Configuration

Virtual Chassis Overview

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.

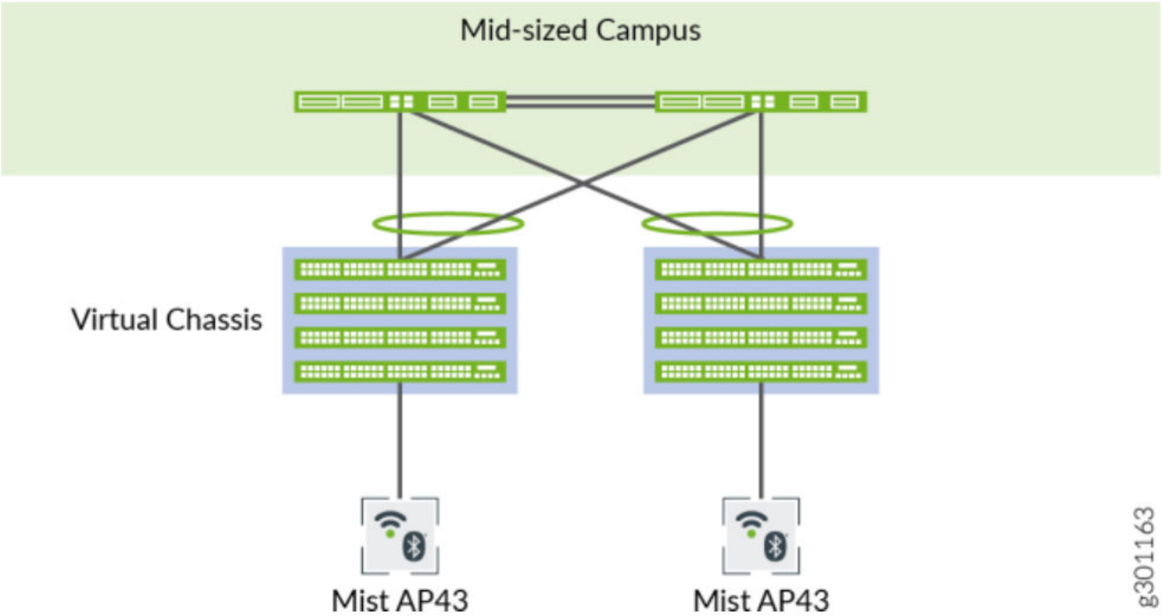

A Virtual Chassis helps you mitigate the risk of loops. It also eliminates the need for legacy redundancy protocols such as spanning tree protocols (STPs) and Virtual Router Redundancy Protocol (VRRP). In core and distribution deployments, you can connect to the Virtual Chassis using link aggregation 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, EX4400, and EX4600 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, EX4300, EX4400, or EX4600 Switches" on page 95 . However, the Modify Virtual Chassis workflow (see "Modify a Virtual Chassis" on page 100) 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 5: Maximum Number of Member Switches Allowed in a Virtual Chassis Configuration

Switch Model	Maximum Member Switches Allowed
EX2300	4
EX4650	4
EX3400	10
EX4100	10
EX4100-F	10
EX4300	10
EX4400	10
EX4600	10
QFX5120-32C, QFX5120-48T, QFX5120-48Y	2
QFX5120-48YM	4

Table 5: Maximum Number of Member Switches Allowed in a Virtual Chassis Configuration *(Continued)*

Switch Model	Maximum Member Switches Allowed
QFX5110	10

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 6: 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)

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 3: Virtual Chassis Setup in a NOC

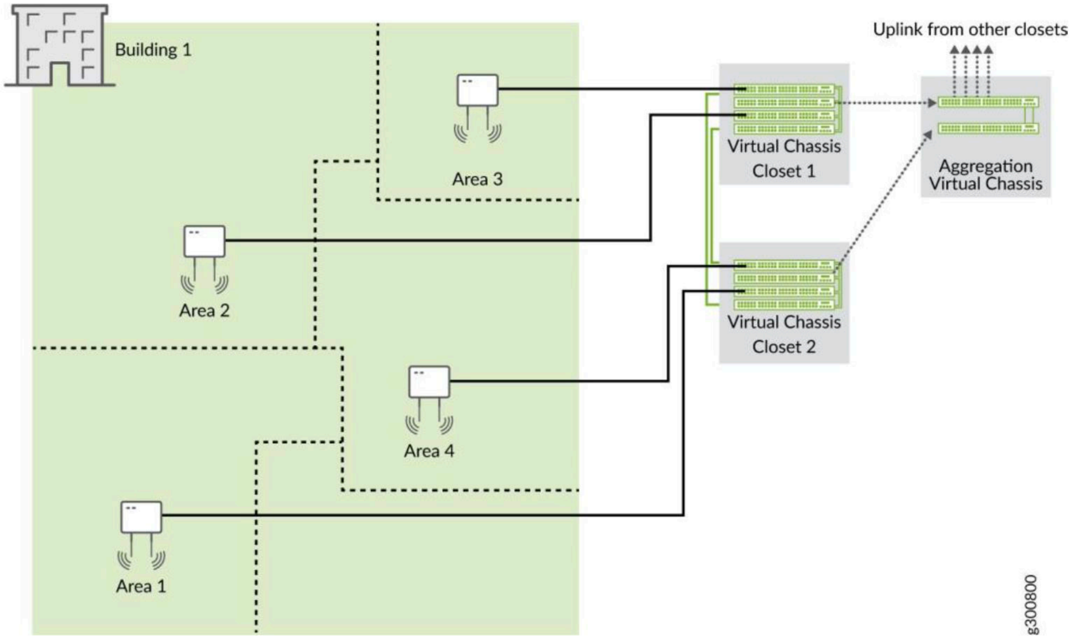

For example, you can use one of the following two options if you want to deploy a solution that includes 96 ports:

- 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 EX2300, EX4650, or QFX5120 Switches

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, EX4400, or EX4600 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, EX4300, EX4400, or EX4600 Switches" on page 95 .

Before configuring a Virtual Chassis, ensure the following:

- 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) to Juniper Mist, see "Onboard a Brownfield Switch" on page 19 . Navigate to the Switches page on the Mist portal and verify that all the switches that you onboarded are listed on the page.

- The switches are connected to the Mist cloud and have the configuration management option enabled on the Mist portal.

  > **NOTE**: The switches need to have a direct connection to the Mist cloud. Ensure that you have an uplink connection directly to the switch.

- All the switches are running the same Junos software version. If they are not, you can upgrade the switch software using a USB drive locally or using the Juniper Mist portal. See "Upgrade Junos OS Software on Your Switch" on page 66 .

To configure a Virtual Chassis using EX2300, EX4650, or QFX5120 switches:

1. Click the **Switches** tab on the left to navigate to the Switches page.
2. 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.

3. Click **More** > **Form Virtual Chassis**.

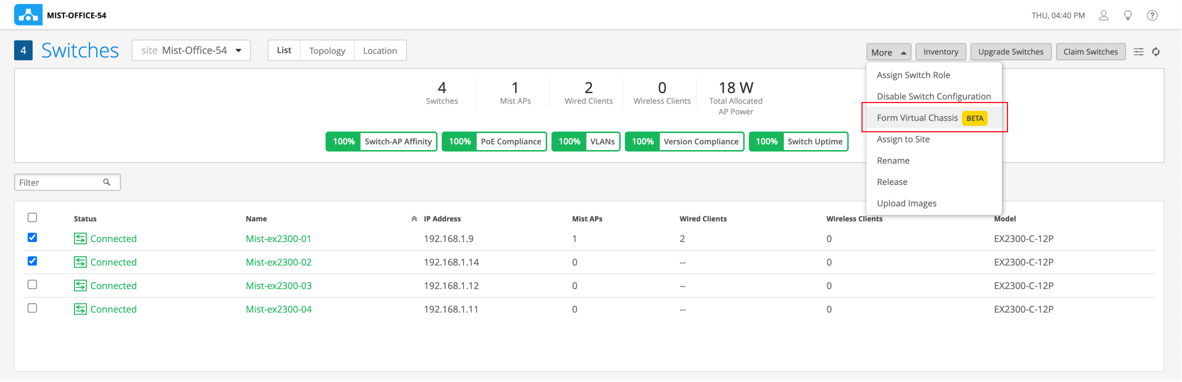

☐	Status	Name	IP Address	Mist APs	Wired Clients	Wireless Clients	Model
☑	Connected	Mist-ex2300-01	192.168.1.9	1	2	0	EX2300-C-12P
☑	Connected	Mist-ex2300-02	192.168.1.14	0	--	0	EX2300-C-12P
☐	Connected	Mist-ex2300-03	192.168.1.12	0	--	0	EX2300-C-12P
☐	Connected	Mist-ex2300-04	192.168.1.11	0	--	0	EX2300-C-12P

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 sample picture.

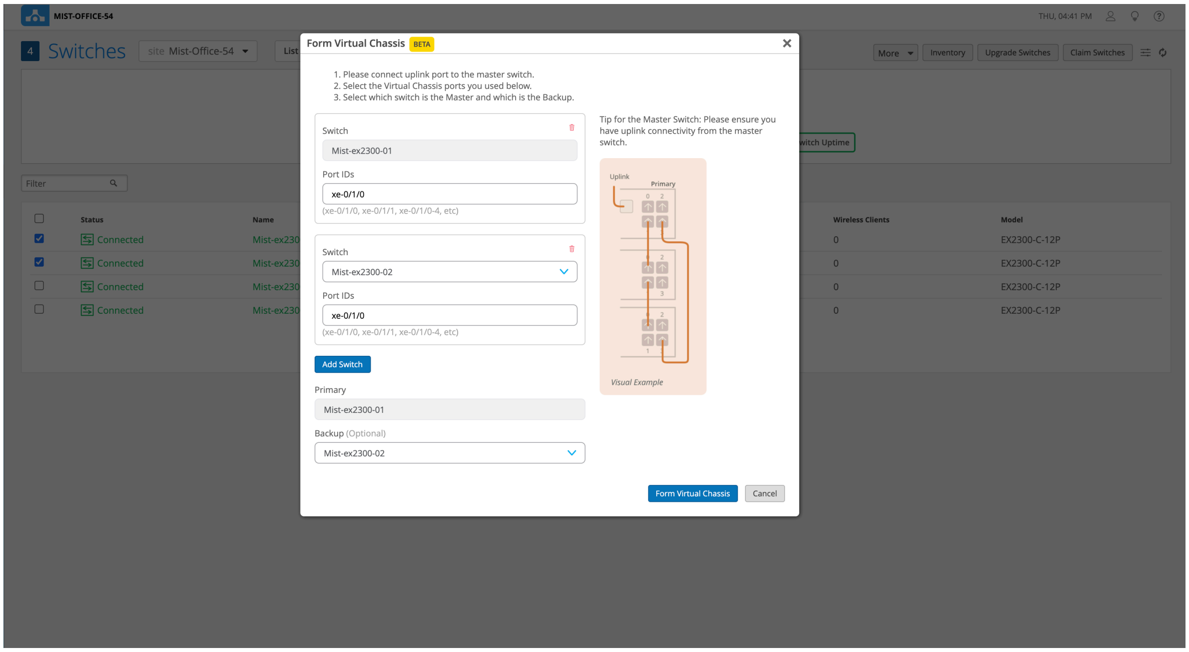

Form Virtual Chassis BETA

1. Please connect uplink port to the master switch.
2. Select the Virtual Chassis ports you used below.
3. Select which switch is the Master and which is the Backup.

Switch: Mist-ex2300-01
Port IDs: xe-0/1/0 (xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

Switch: Mist-ex2300-02
Port IDs: xe-0/1/0 (xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

Primary: Mist-ex2300-01

Backup (Optional): Mist-ex2300-02

Tip for the Master Switch: Please ensure you have uplink connectivity from the master switch.

4. 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 **Primary** switch. The switch that you selected first is the primary switch by default. You can modify that.

   c. The **Backup** switch. This configuration is optional. If you don't select a switch to function in the backup role, Mist assigns the linecard role to that switch.

   Ensure that you have an uplink connection directly to the primary switch.

5. Click **Form Virtual Chassis** and wait for 3 to 5 minutes 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.

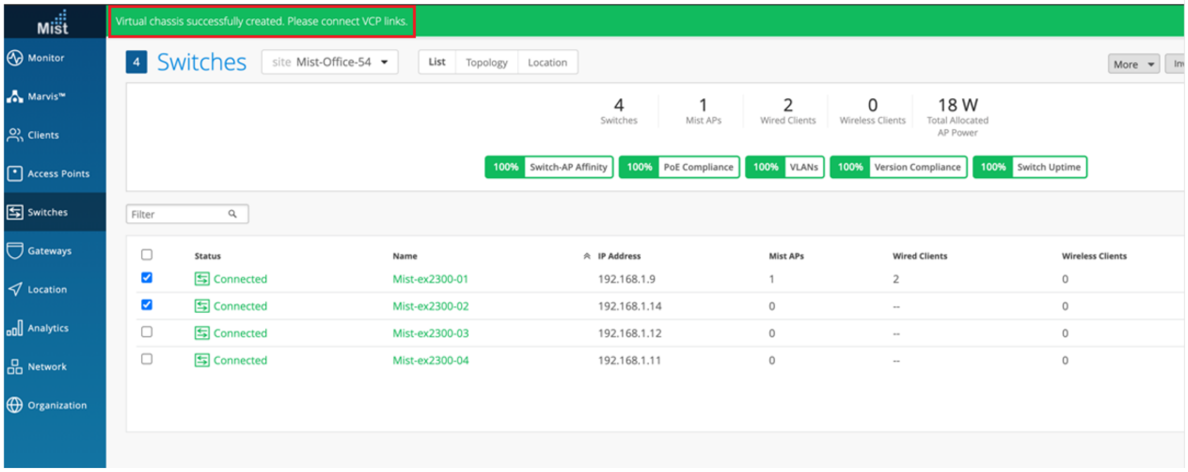

Virtual chassis successfully created. Please connect VCP links.

	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

6. Connect the switches to each other using the VCPs configured.

   When the Virtual Chassis formation is in progress, the Switches page shows the switch status as **VC forming**.

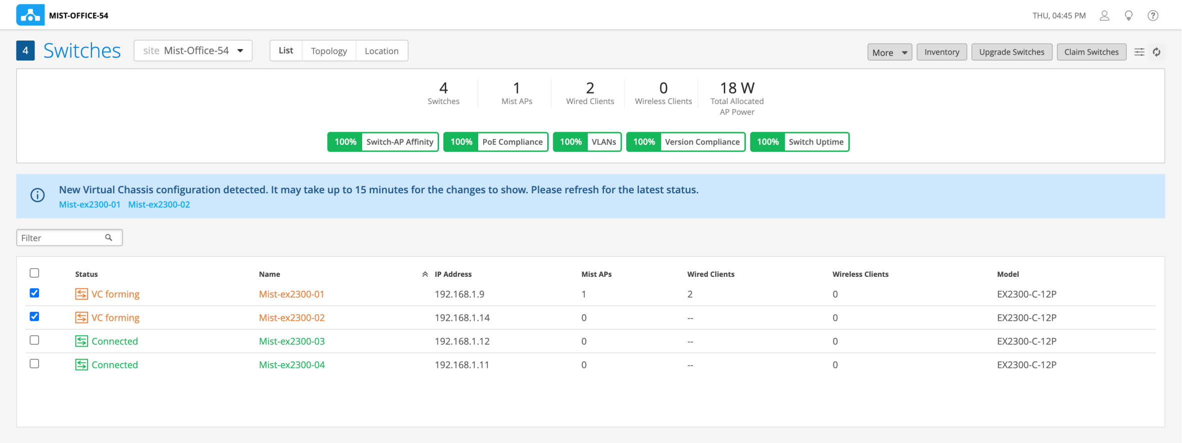

MIST-OFFICE-54 THU, 04:45 PM

4 Switches site Mist-Office-54 List Topology Location More Inventory Upgrade Switches Claim Switches

4 Switches 1 Mist APs 2 Wired Clients 0 Wireless Clients 18 W Total Allocated AP Power

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

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

Filter

□	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.

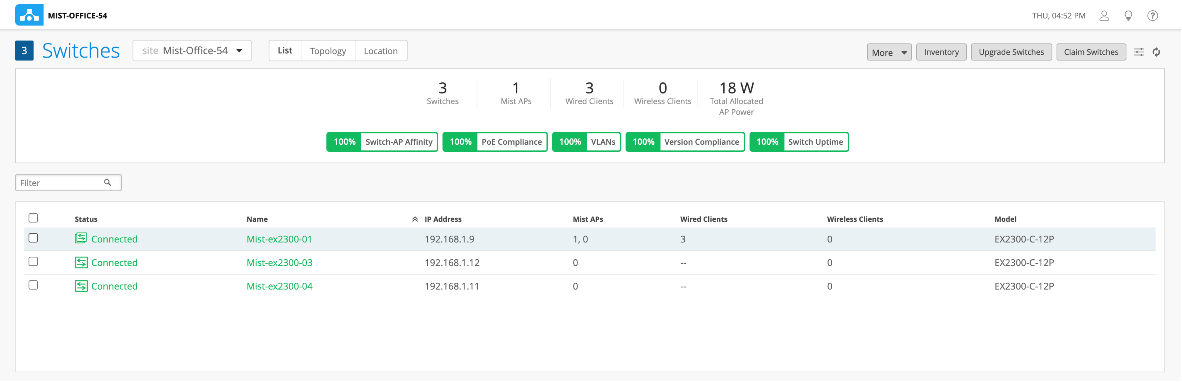

MIST-OFFICE-54 THU, 04:52 PM

3 Switches site Mist-Office-54 List Topology Location More Inventory Upgrade Switches Claim Switches

3 Switches 1 Mist APs 3 Wired Clients 0 Wireless Clients 18 W Total Allocated AP Power

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

Filter

□	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.

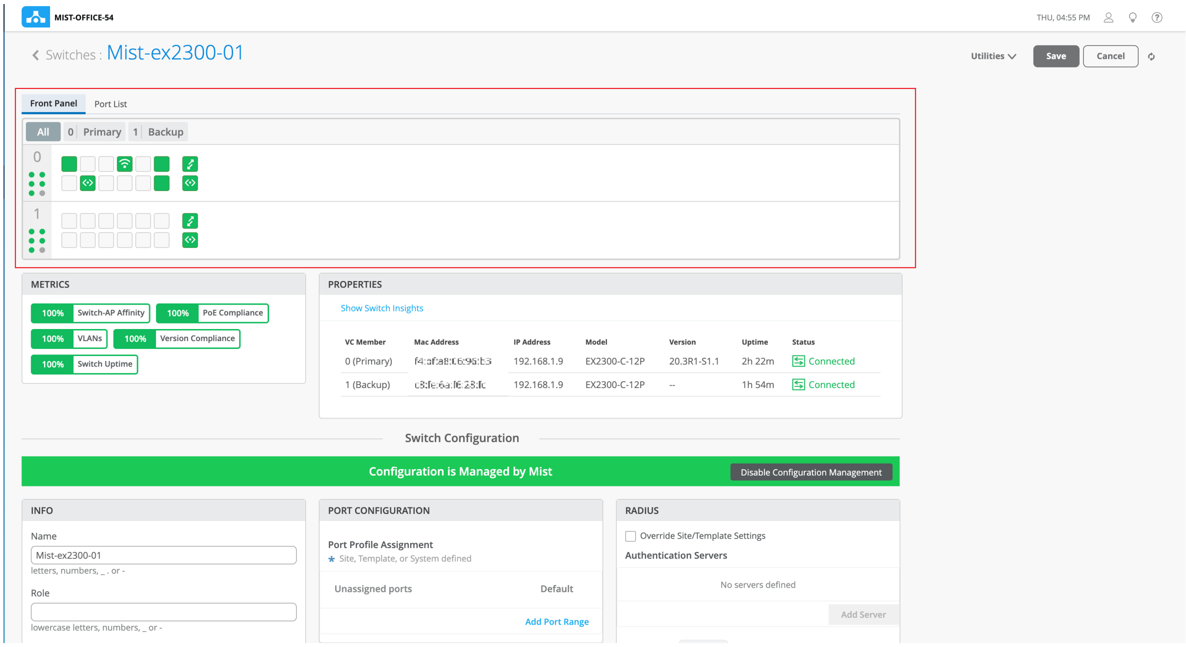

VC Member	Mac Address	IP Address	Model	Version	Uptime	Status
0 (Primary)	[redacted]	192.168.1.9	EX2300-C-12P	20.3R1-S1.1	2h 22m	Connected
1 (Backup)	[redacted]	192.168.1.9	EX2300-C-12P	--	1h 54m	Connected

Name: Mist-ex2300-01
letters, numbers, _ . or -

Role:
lowercase letters, numbers, _ or -

Unassigned ports	Default

□ Override Site/Template Settings

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 "Modify a Virtual Chassis" on page 100 .

Configure a Virtual Chassis Using EX3400, EX4100, EX4100-F, EX4300, EX4400, or EX4600 Switches

The EX3400, EX4100, EX4100-F, EX4300, EX4400, and EX4600 switches come with dedicated Virtual Chassis ports (VCPs). You can create Virtual Chassis using these switches by connecting them to each other via VCPs. These switches don't support the Form Virtual Chassis option on the Switches page on the Mist portal. 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.

Before configuring a Virtual Chassis using the EX3400, EX4100, EX4100-F, EX4300, EX4400, or EX4600 Switches, ensure the following:

- 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 19 .
- All the switches are running the same Junos software version. If they are not, you can upgrade the switch software using a USB drive locally or using the Juniper Mist portal. See "Upgrade Junos OS Software on Your Switch" on page 66 .

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. The Modify Virtual Chassis option is available for switches that have the configuration management enabled in Mist.

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 89 .

Table 7: 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)

To configure a Virtual Chassis using EX3400, EX4100, EX4100-F, EX4300, EX4400, or EX4600 switches:

1. Power off the switches that you want to include in the Virtual Chassis.
2. 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.

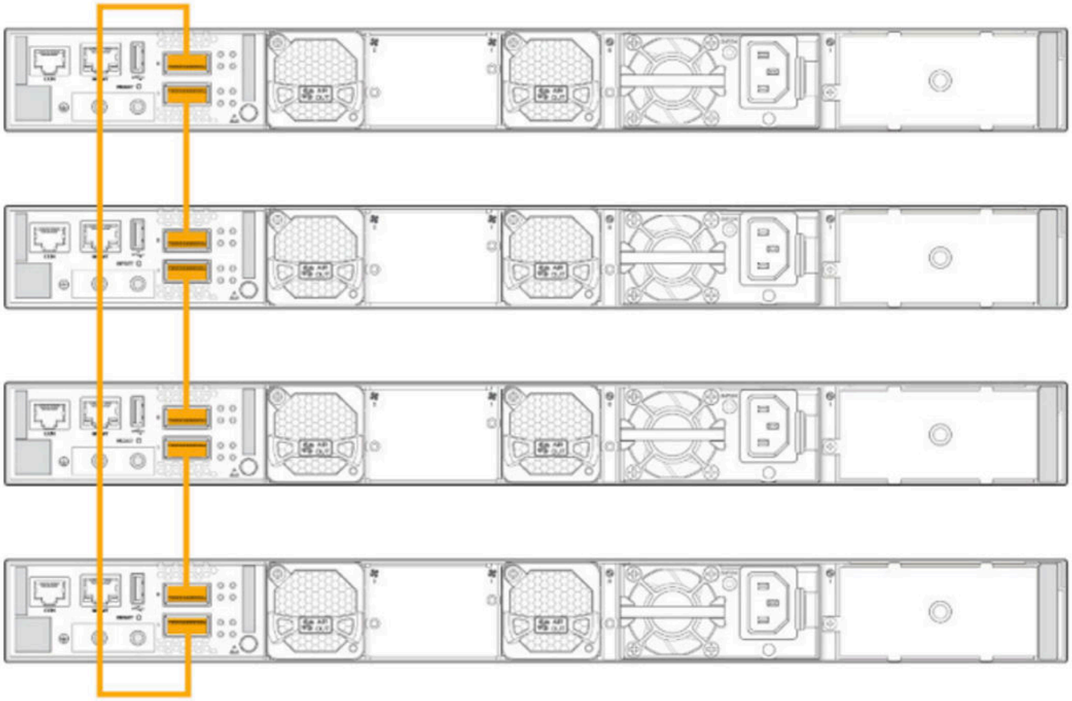

3. Power on the switch that you want to function in the primary role.

   This member will become FPC0.

   > **NOTE**: The order in which you power on devices also determines the member ID. 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 power them on and then connect them to the other existing switches in that order.

4. Approximately one minute after powering on the switch that you selected for the primary role, power on the switch that you want to function in the backup role.

   > **NOTE**: In the case of some switch models, you may need to wait for more than one minute as they take more time to boot up.

   This member will become FPC1.

5. Wait for approximately one more minute, and then boot up the rest of the switches that you want to function in the linecard role.

NOTE: In the case of some switch models, you may need to wait for more than one minute as they take more time to boot up.

6. 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.
   A Virtual Chassis is now physically formed but not preprovisioned.
7. Connect the Virtual Chassis to the Juniper Mist cloud by connecting the uplink port on the primary switch to the upstream switch.
   This step initiates a zero-touch provisioning (ZTP) process on the Virtual Chassis and connects it to the Juniper Mist cloud.
8. 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:

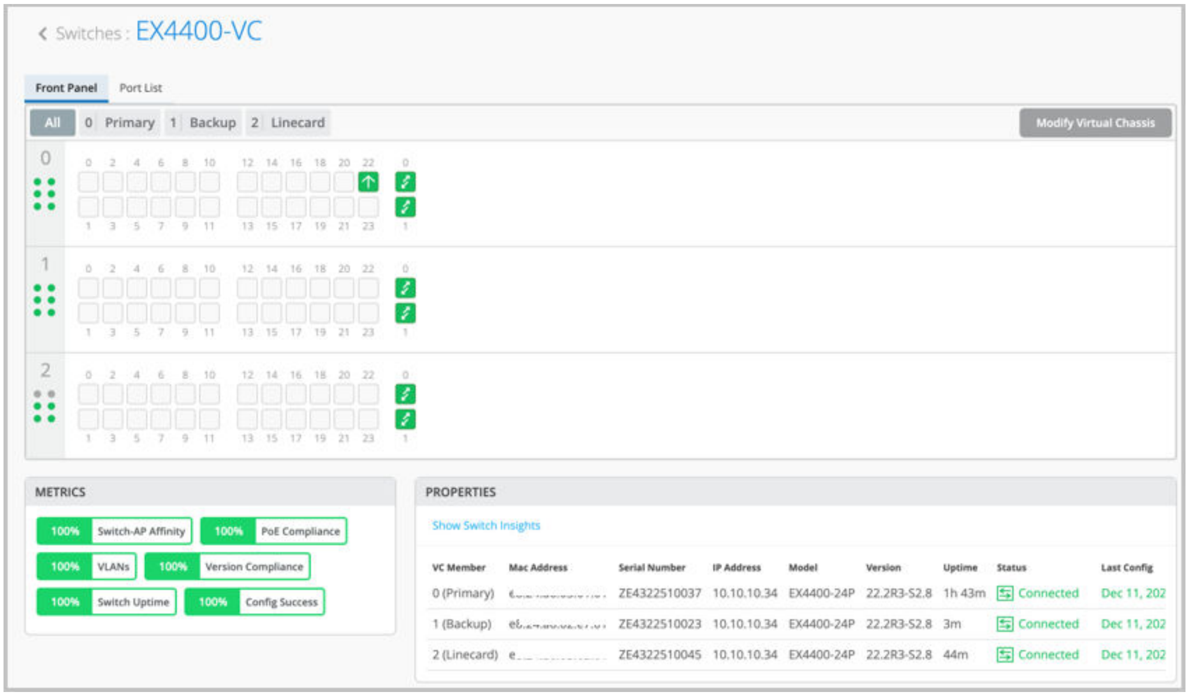

100%	Switch-AP Affinity
100%	PoE Compliance
100%	VLANs
100%	Version Compliance
100%	Switch Uptime
100%	Config Success

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last Config
0 (Primary)	[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

9. 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.

   b. On the Modify Virtual Chassis page, click **Preprovision Virtual Chassis**. See a sample Modify Virtual Chassis page below:

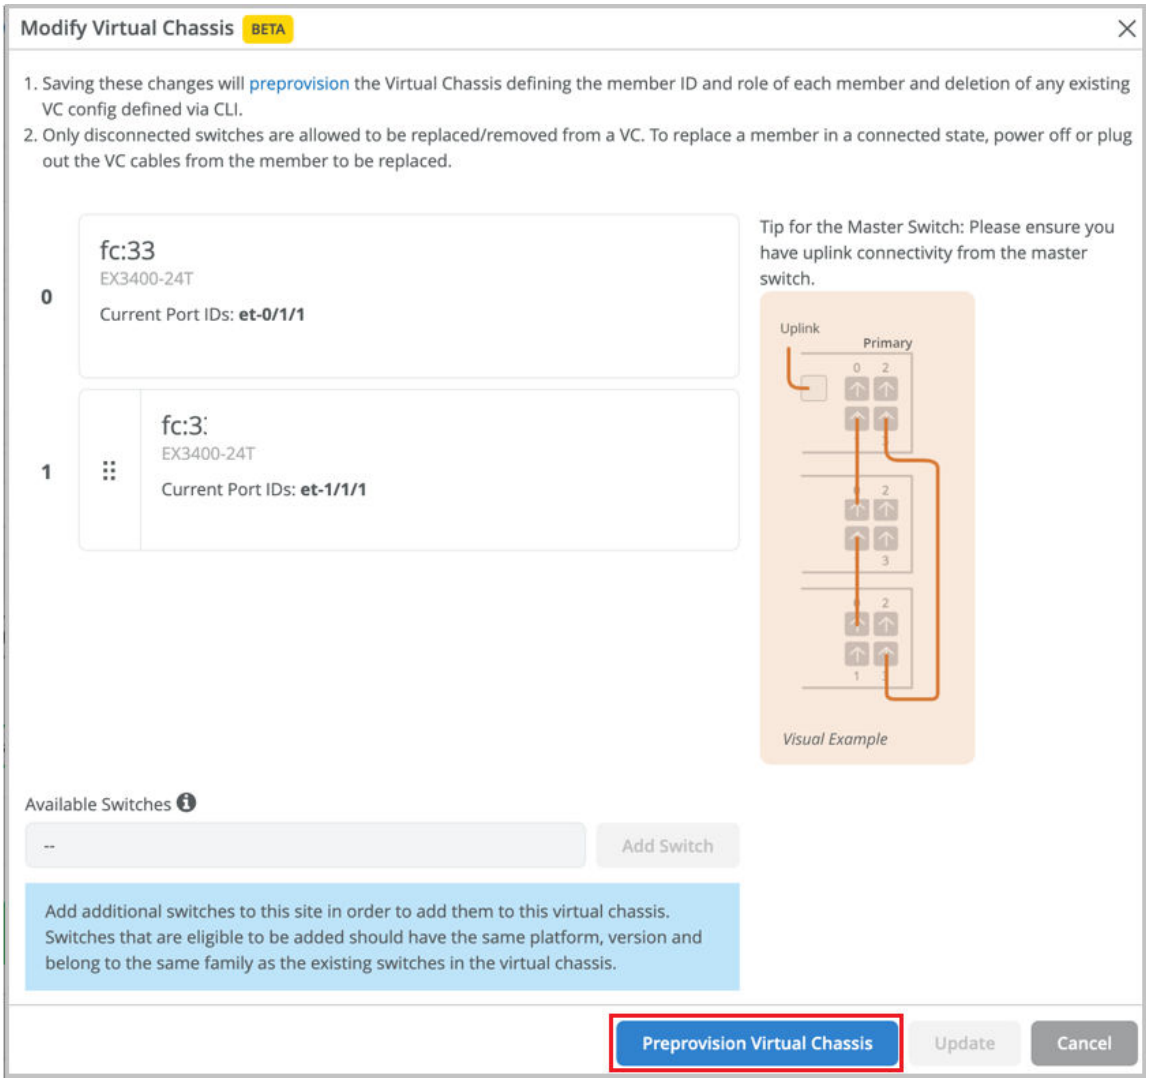

Modify Virtual Chassis BETA

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 fc:33
EX3400-24T
Current Port IDs: **et-0/1/1**

1 fc:3:
EX3400-24T
Current Port IDs: **et-1/1/1**

Tip for the Master Switch: Please ensure you have uplink connectivity from the master 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.

This step pushes the preprovisioned Virtual Chassis configuration to the device and overwrites the old autoprovision Virtual Chassis configuration in the device. 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 "Modify a Virtual Chassis" on page 100 .

Modify a Virtual Chassis

IN THIS SECTION

You can use the **Modify Virtual Chassis** option on the switch details page to renumber and replace the Virtual Chassis members and add new members to a Virtual Chassis.

The Modify Virtual Chassis workflow leverages the pre-provisioned way of Junos configuration. This option is visible for switches that have the configuration management option 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.

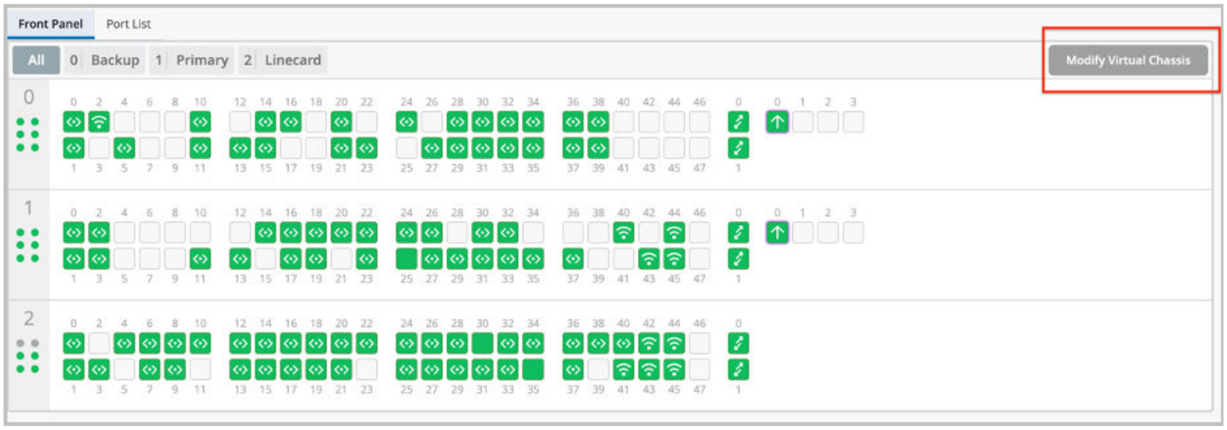

NOTE:

- 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 the FPC0, trash and replace it with an existing member in the Virtual Chassis. You cannot add a new member during the deletion of the FPC0.
- The Add Switch dropdown only shows the switches that:
  - Share the same major firmware version as the existing members in the Virtual Chassis.
  - 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.
- 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.

The Modify Virtual Chassis workflow leverages the Junos preprovisioning method which configures the role and serial number of all members in a Virtual Chassis. To learn more about preprovisioning, see Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File.

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.

Replace a Member Switch in a Virtual Chassis

IN THIS SECTION

You can replace a disconnected Virtual Chassis member switch with another, by deleting the old member and adding a new member. For this feature to work, you must ensure the following:

- The new switch is of the same model as the other members in the Virtual Chassis.
- The new switch runs the same Junos version as the other members.
- The new switch is connected to the network.
- The new switch is assigned to the same site as the other members in the Virtual Chassis.

Replace a Non-FPC0 Member in a Virtual Chassis

To replace a non-FPC0 member:

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.

   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 19 .

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. Ensure that the new switch is running the same Junos version as the other members in the Virtual Chassis. If it is not, upgrade the switch using a USB drive locally or using the Juniper Mist portal. See "Upgrade Junos OS Software on Your Switch" on page 66 for more information.
3. Power off the new switch (the replacement switch). If you are replacing the FPC0 member, you must keep the replacement switch powered on all the time.
4. Power off the member to be replaced. Or, remove the Virtual Chassis port (VCP) cables from this member.
5. Connect the Virtual Chassis cables from the existing Virtual Chassis members to the new replacement switch.
6. On the Mist portal, navigate to the switch details page of the Virtual Chassis by clicking **Switches** > *Switch Name* (the Virtual Chassis name).
7. Wait for the switch details page (the Virtual Chassis page) to display the member switch to be replaced as offline, as shown below:

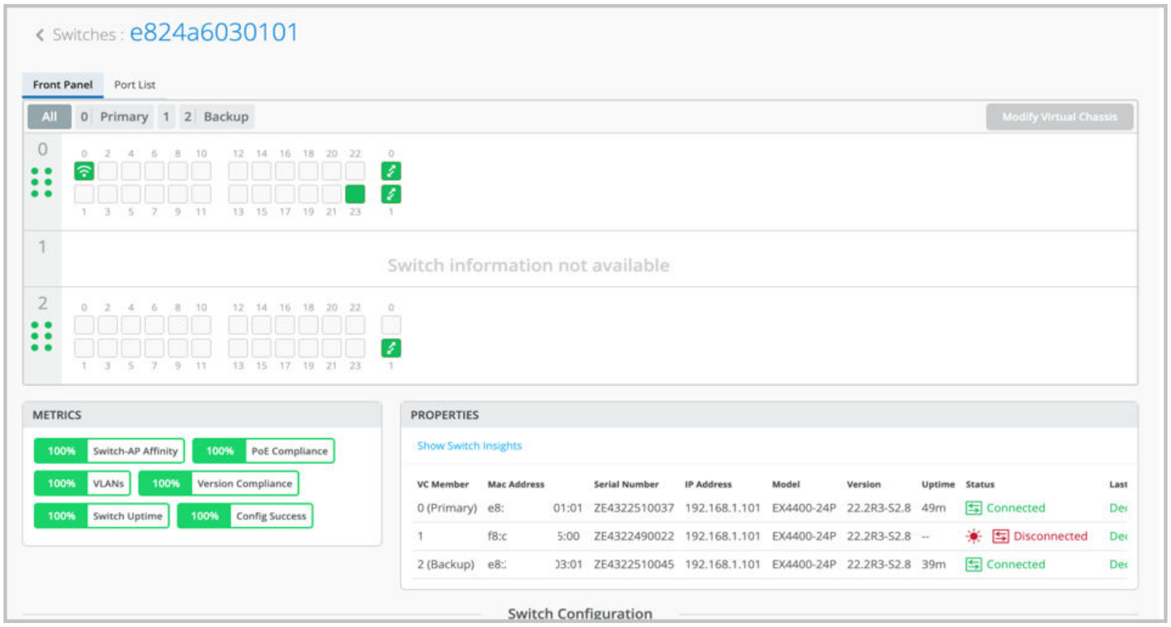

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last
0 (Primary)	e8: 01:01	ZE4322510037	192.168.1.101	EX4400-24P	22.2R3-S2.8	49m	Connected	De(
1	f8:c 5:00	ZE4322490022	192.168.1.101	EX4400-24P	22.2R3-S2.8	--	Disconnected	De(
2 (Backup)	e8::)3:01	ZE4322510045	192.168.1.101	EX4400-24P	22.2R3-S2.8	39m	Connected	De(

8. Click **Modify Virtual Chassis**.

   Because you have 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.

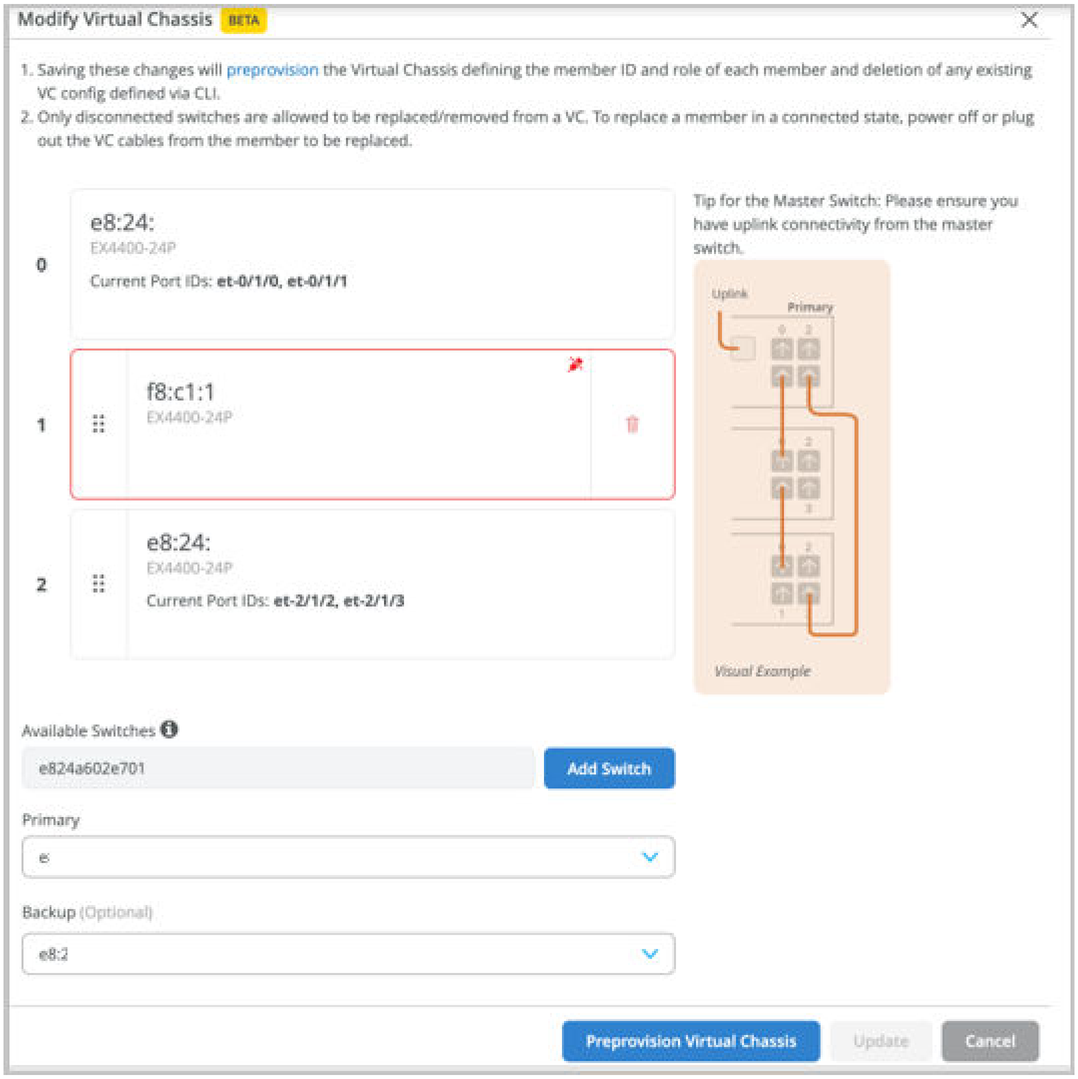

Modify Virtual Chassis BETA

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:
EX4400-24P
Current Port IDs: **et-0/1/0, et-0/1/1**

1 f8:c1:1
EX4400-24P

2 e8:24:
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
e824a602e701

Primary
e

Backup (Optional)
e8:2

9. Delete the member to be replaced by clicking the trash icon.
10. Click **Add Switch** to add the new replacement member.
11. Renumber the new switch by dragging and dropping it into the appropriate slot. Remember to edit the MAC address of the backup switch if you are replacing the backup switch.
12. Click **Update**.

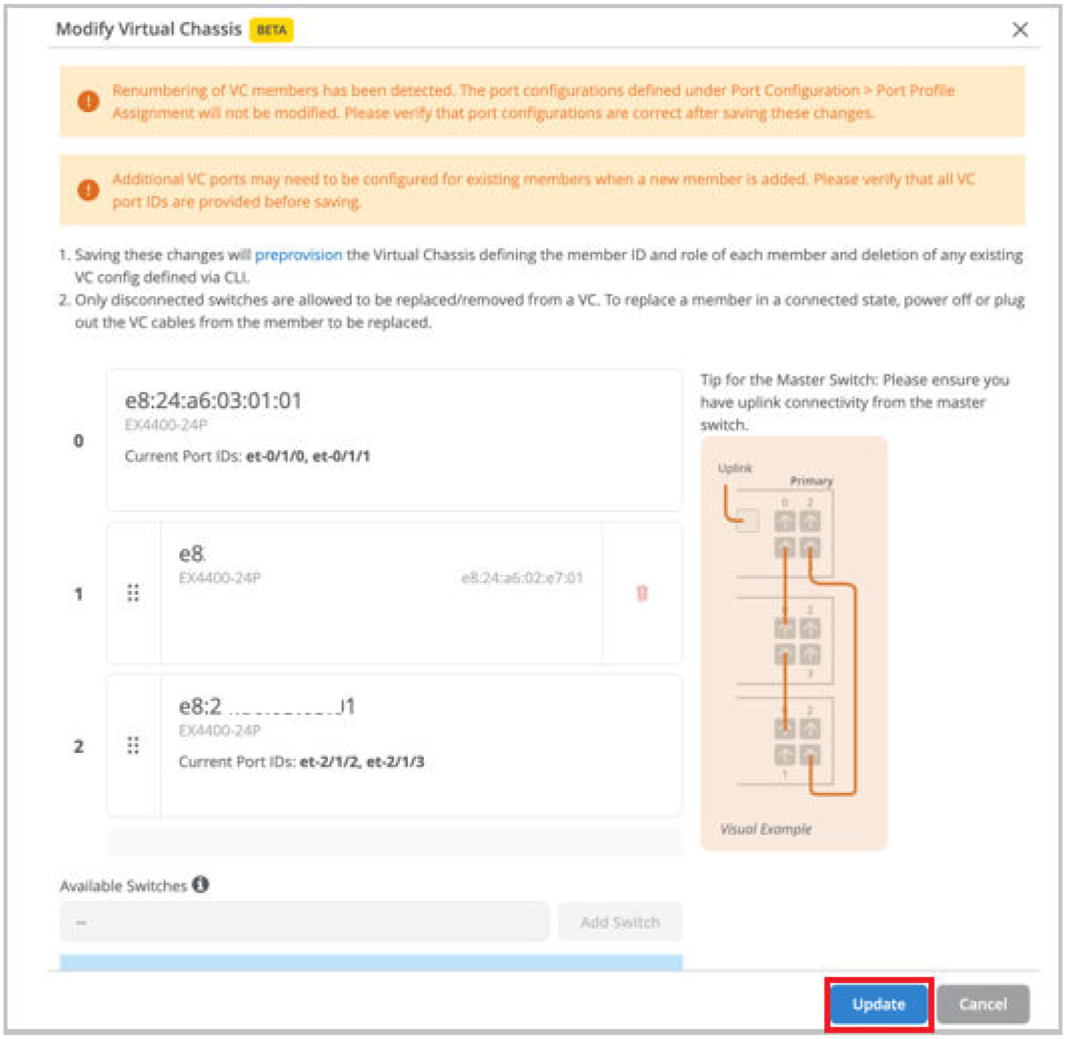

Modify Virtual Chassis BETA

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 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. Power on the replacement switch and wait for Virtual Chassis formation to be complete.

    The Switch Events page displays all the Virtual Chassis update events.

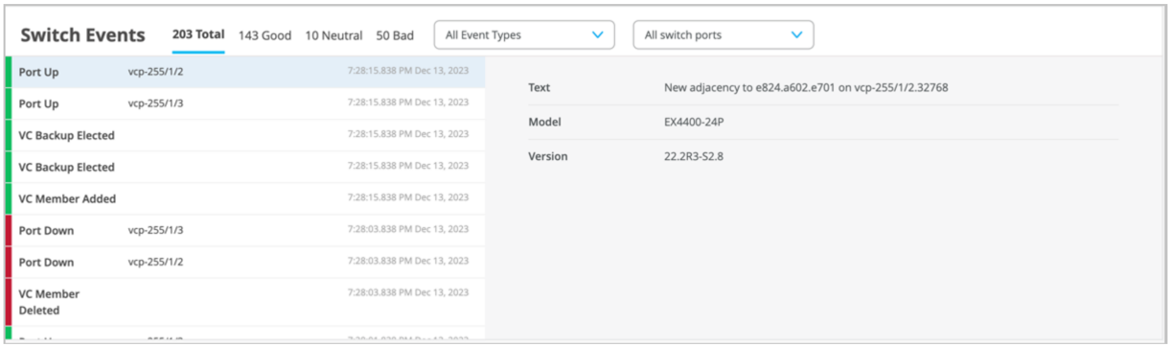

Switch Events 203 Total 143 Good 10 Neutral 50 Bad

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.

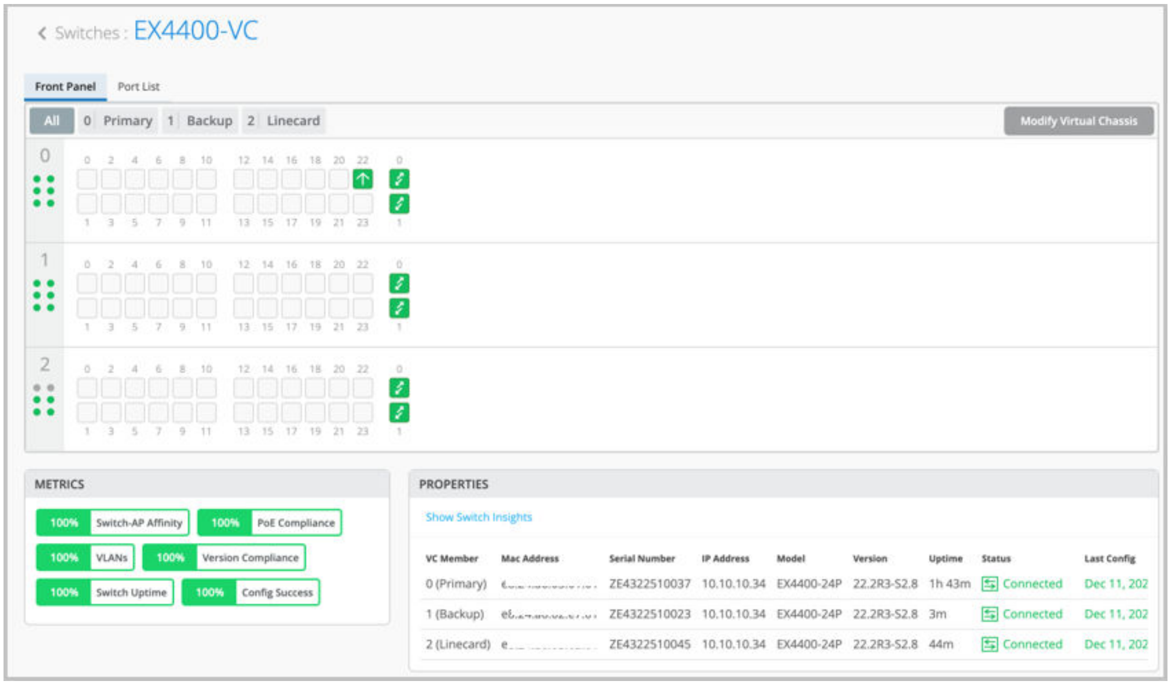

100% Switch-AP Affinity
100% PoE Compliance
100% VLANs
100% Version Compliance
100% Switch Uptime
100% Config Success

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last Config
0 (Primary)	[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 the FPC0 Member in a Virtual Chassis

The FPC0 is used as the device identifier and is used to communicate to the Mist cloud. You cannot replace the FPC0 with another member in a single operation. You need to follow a 2-step process - adding the new replacement switch and then removing the switch to be replaced. You should carry out the FPC0 replacement operation in a maintenance window as this operation can impact the traffic to the clients connected.

1. Click **Switches** > *Switch Name* to go to the switch details page of the Virtual Chassis to be modified.
2. Click **Modify Virtual Chassis**.

   The Modify Virtual Chassis window appears.
3. On the Modify Virtual Chassis window, click **Add Switch** and add the replacement switch to the Virtual Chassis as a new member.

   The new switch must be powered on and connected to the Virtual Chassis.

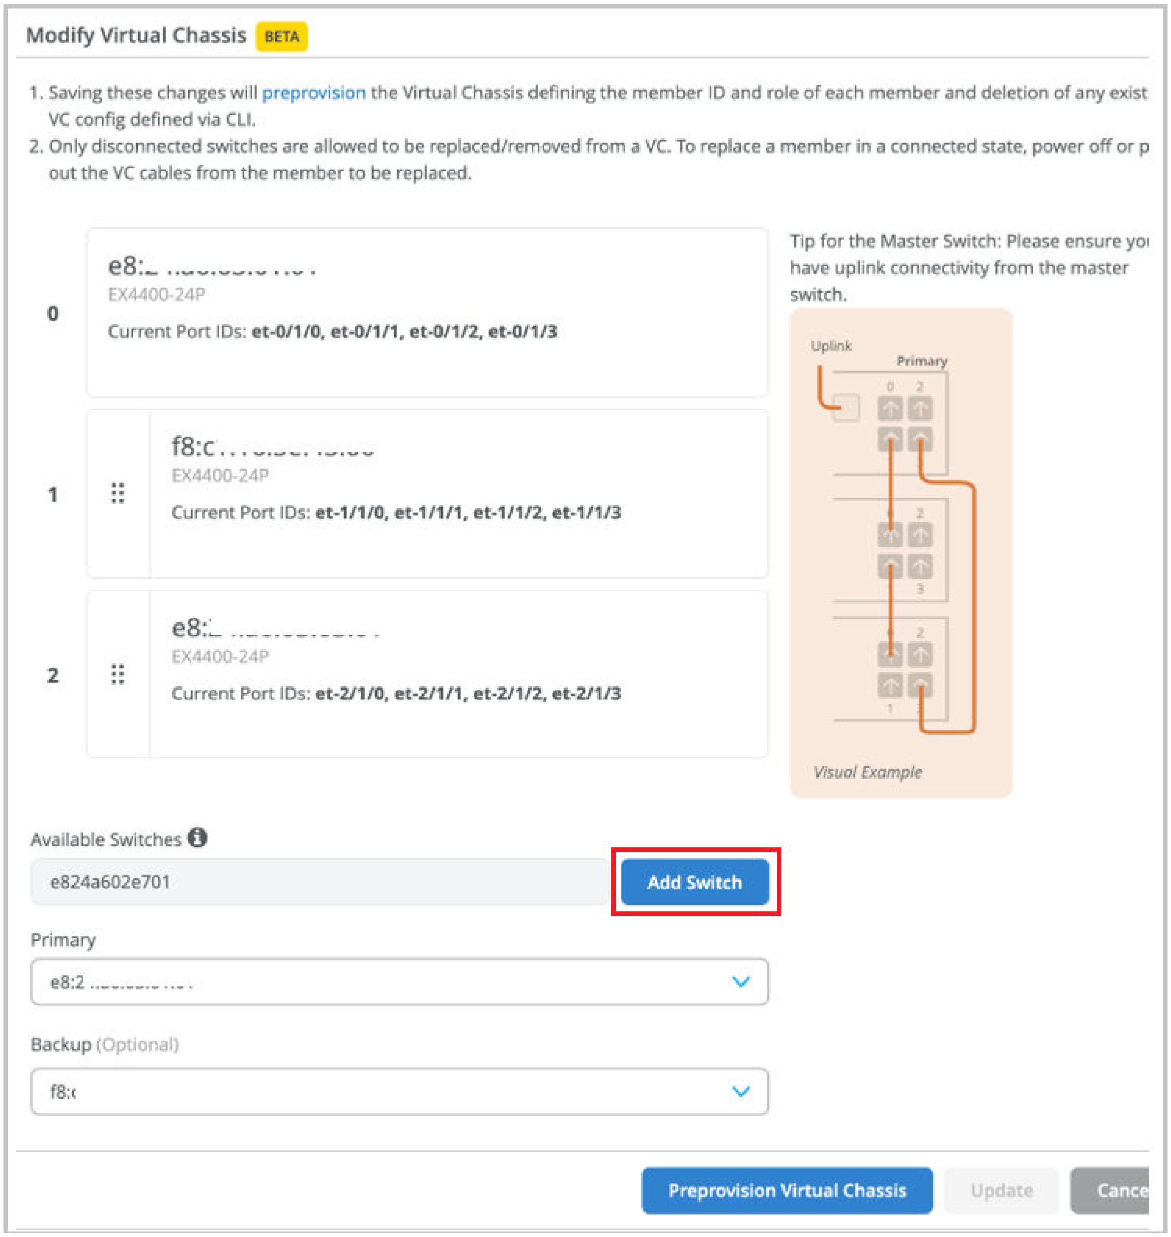

Modify Virtual Chassis BETA

1. Saving these changes will preprovision the Virtual Chassis defining the member ID and role of each member and deletion of any exist 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 p out the VC cables from the member to be replaced.

0 e8:[redacted]
EX4400-24P
Current Port IDs: **et-0/1/0, et-0/1/1, et-0/1/2, et-0/1/3**

1 f8:c[redacted]
EX4400-24P
Current Port IDs: **et-1/1/0, et-1/1/1, et-1/1/2, et-1/1/3**

2 e8:[redacted]
EX4400-24P
Current Port IDs: **et-2/1/0, et-2/1/1, et-2/1/2, et-2/1/3**

Tip for the Master Switch: Please ensure yo have uplink connectivity from the master switch.

Available Switches: e8824a602e701

Primary: e8:2[redacted]

Backup (Optional): f8:[redacted]

4. Click **Update**.

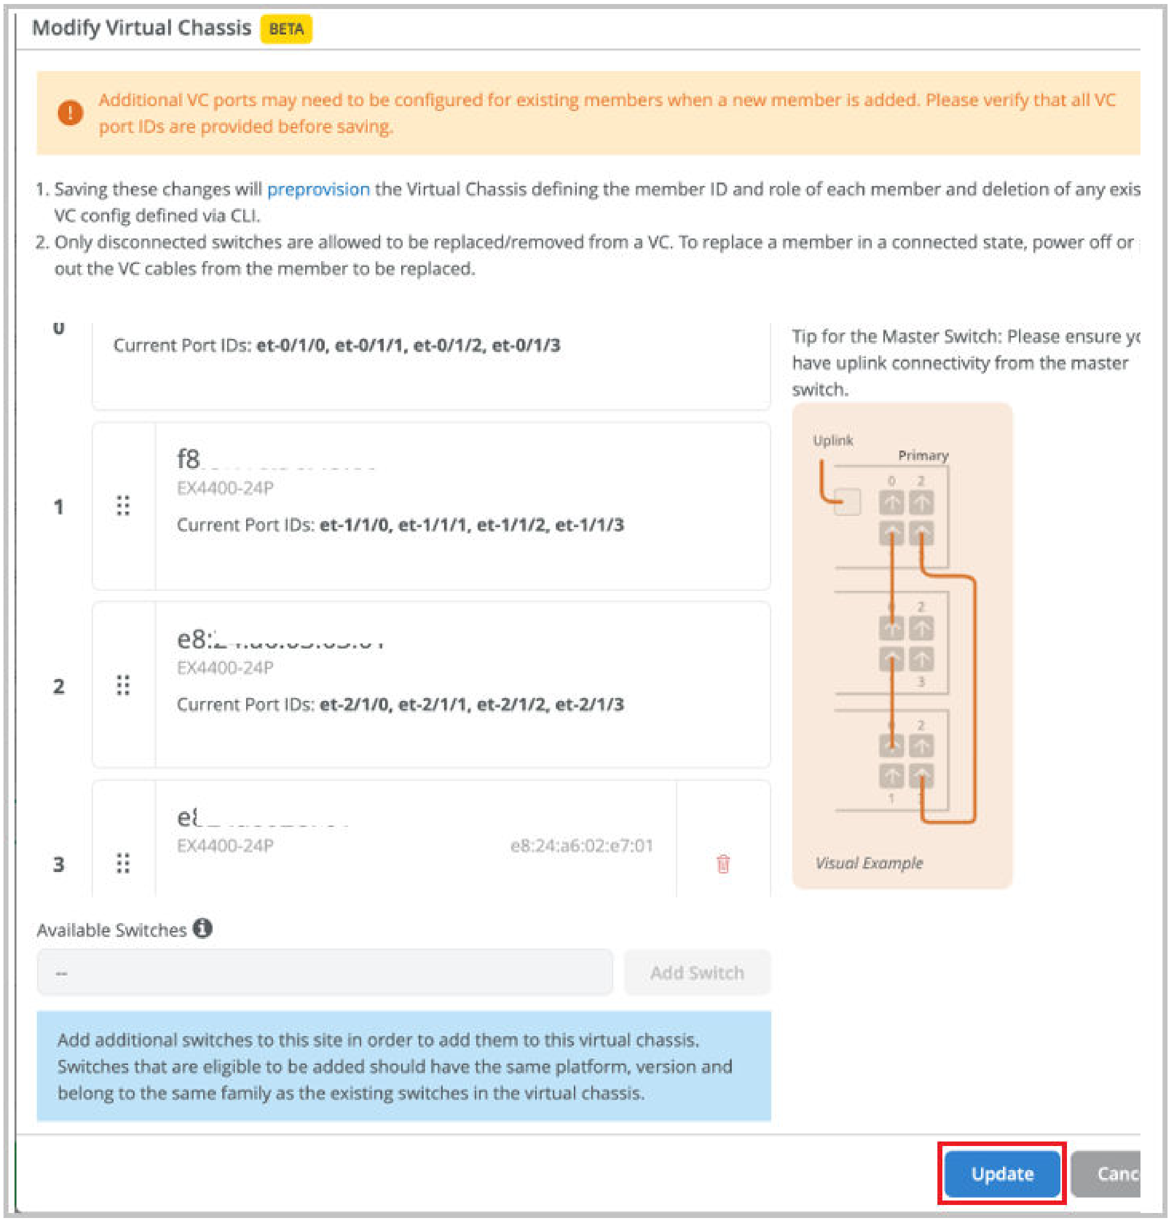

Additional VC ports may need to be configured for existing members when a new member is added. Please verify that all VC port IDs are provided before saving.

1. Saving these changes will preprovision the Virtual Chassis defining the member ID and role of each member and deletion of any exis 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 out the VC cables from the member to be replaced.

0 Current Port IDs: **et-0/1/0, et-0/1/1, et-0/1/2, et-0/1/3**

1 f8[redacted]
EX4400-24P
Current Port IDs: **et-1/1/0, et-1/1/1, et-1/1/2, et-1/1/3**

2 e8:[redacted]
EX4400-24P
Current Port IDs: **et-2/1/0, et-2/1/1, et-2/1/2, et-2/1/3**

3 e[redacted]
EX4400-24P e8:24:a6:02:e7:01

Tip for the Master Switch: Please ensure yo have uplink connectivity from the master 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.

5. Remove the uplink connection (in-band or OOB) from 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.
6. Power off the FPC0 member to be replaced. Or, remove the VCP cables from it.
7. Remove the DAC cables from the FPC0 being replaced and connect it to the new member in the same ports.

   The Mist cloud adds the new member to the Virtual Chassis.
8. On the Mist portal, navigate to the switch details page of the Virtual Chassis by clicking **Switches** > *Switch Name* (the Virtual Chassis name).
   The switch details page (the Virtual Chassis page) will display the new member switch as part of the Virtual Chassis.

9. On the switch details page, click **Modify Virtual Chassis**.

   Because you have 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.

10. Delete the member to be replaced by clicking the trash icon.
    The Modify Virtual Chassis window displays a message indicating that FPC0 is required.

11. Move the new switch to slot 0 (the FPC0 slot) by dragging and dropping it. Also, update the **Backup** field with the MAC address of the new switch, as shown below:

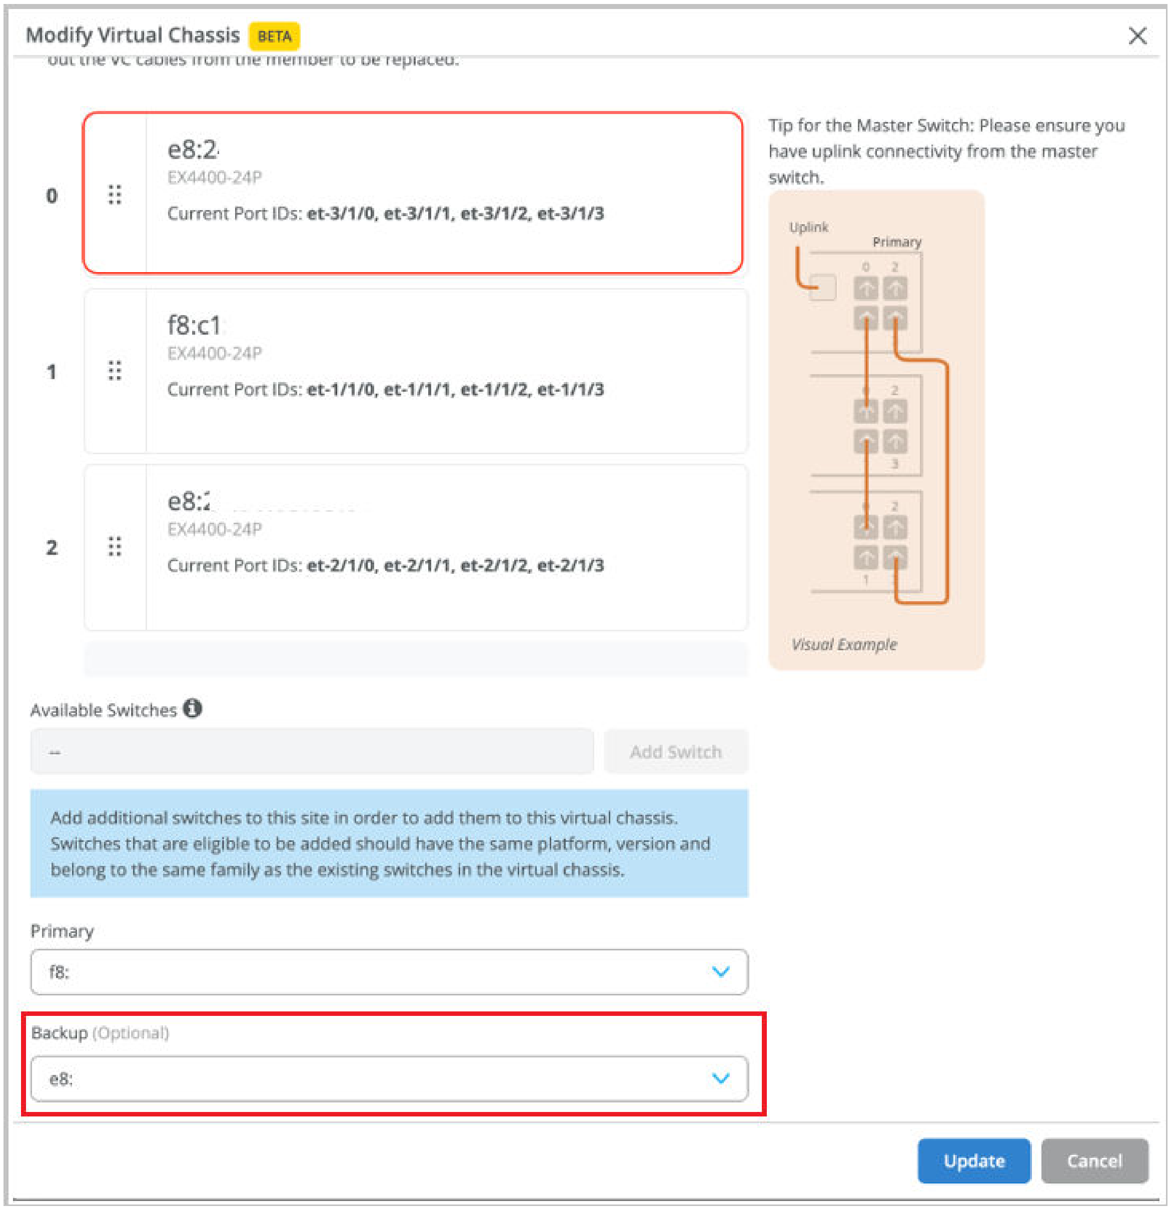

Modify Virtual Chassis BETA

out the VC cables from the member to be replaced.

0 e8:2
EX4400-24P
Current Port IDs: **et-3/1/0, et-3/1/1, et-3/1/2, et-3/1/3**

1 f8:c1
EX4400-24P
Current Port IDs: **et-1/1/0, et-1/1/1, et-1/1/2, et-1/1/3**

2 e8:
EX4400-24P
Current Port IDs: **et-2/1/0, et-2/1/1, 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

--

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.

Primary: f8:

Backup (Optional): e8:

This step links the new FPC0 device MAC address to the web browser URL and updates the outbound SSH MAC address field with the new FPC0 device MAC address.

12. 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 power these members on and then connect them to the other existing Virtual Chassis switch members in that order.

You can modify the member switches' order on the Mist portal by renumbering the members. On the Modify Virtual Chassis window, accessed 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.

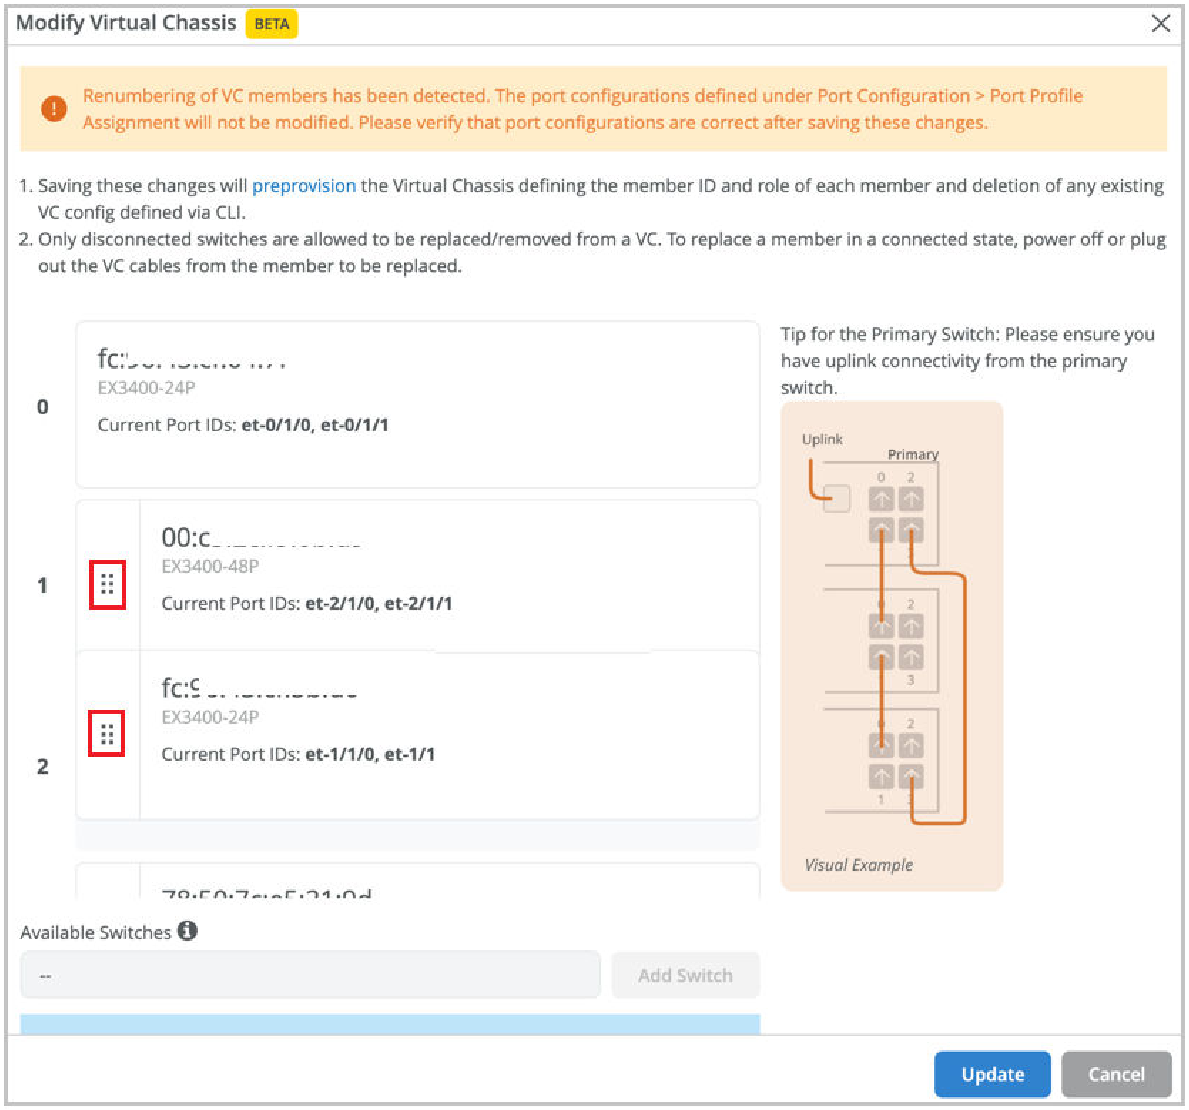

Modify Virtual Chassis BETA

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 fc:[redacted]
EX3400-24P
Current Port IDs: **et-0/1/0, et-0/1/1**

1 00:c[redacted]
EX3400-48P
Current Port IDs: **et-2/1/0, et-2/1/1**

2 fc:[redacted]
EX3400-24P
Current Port IDs: **et-1/1/0, et-1/1**

78:50:7c:e5:21:0d

Tip for the Primary Switch: Please ensure you have uplink connectivity from the primary switch.

NOTE: You must specify the FPC0 in this configuration. Within a Virtual Chassis, you cannot renumber, move around, or delete FPC0 unless it is disconnected. It is the device identifier for connectivity to the Mist cloud.

Renumbering the members within a Virtual Chassis does not renumber the port configurations and port profile assignment. 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 the 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. You can change the role of a switch from primary to backup or backup to linecard or linecard to primary.

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 a primary switch and a backup switch (optional) from the Primary and Backup drop-down list. All the other switches assume a linecard role.

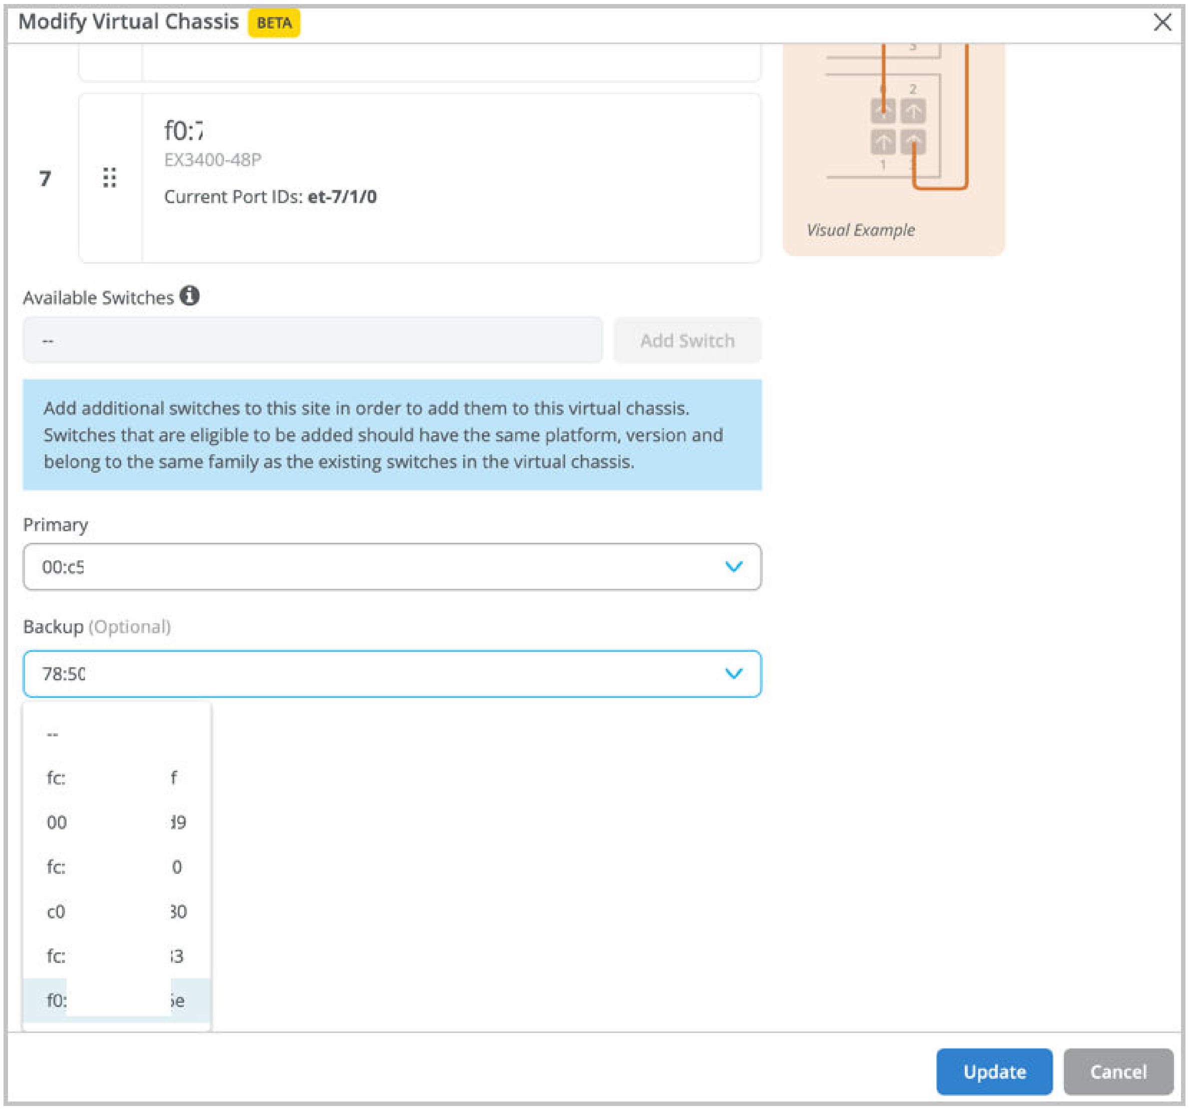

Modify Virtual Chassis BETA

7 f0:7
EX3400-48P
Current Port IDs: **et-7/1/0**

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.

Primary: 00:c5

Backup (Optional): 78:50

--
fc: f
00 19
fc: 0
c0 30
fc: 3
f0: 5e

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:

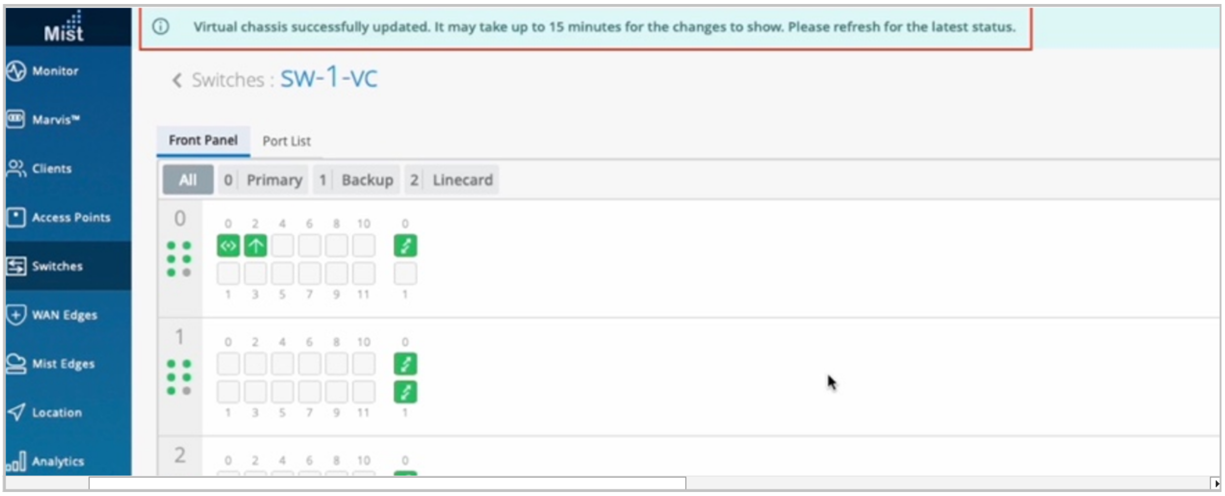

Virtual chassis successfully updated. It may take up to 15 minutes for the changes to show. Please refresh for the latest status.

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. Remove the physical VCP connection of the member switch that you want to delete from the Virtual Chassis.
2. On the Mist portal, click the **Switches** tab on the left to navigate to the Switches page.
3. Click the Virtual Chassis from which you want to delete a member switch.
   The switch details page appears.
4. 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.
5. Click the delete icon and then click **Update**.

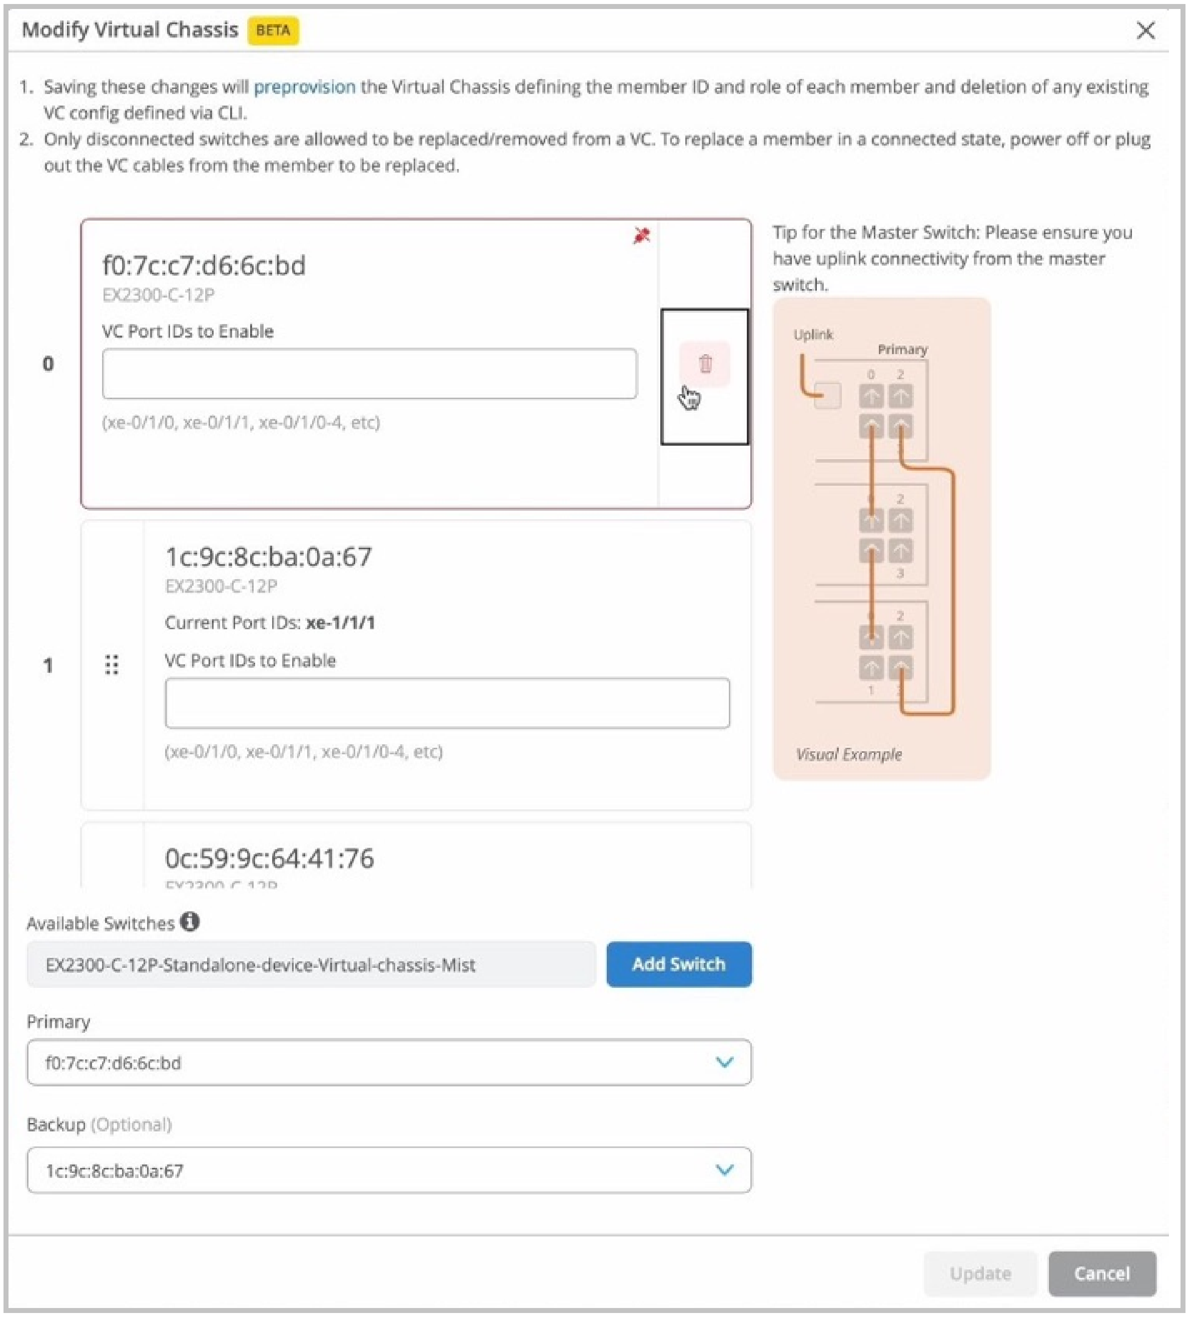

Modify Virtual Chassis BETA

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

Primary: f0:7c:c7:d6:6c:bd

Backup (Optional): 1c:9c:8c:ba:0a:67

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 as the other members in the Virtual Chassis.

- The new switch runs the same Junos version as the other members.
- The new switch is connected to the network.
- The new switch is assigned to the same site as the other members in the Virtual Chassis.

To add a new member switch to the Virtual Chassis:

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.

   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 19 .

   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. Ensure that the new switch is running the same Junos version as the other members in the Virtual Chassis. If it is not, upgrade the switch using a USB drive locally or using the Juniper Mist portal. See "Upgrade Junos OS Software on Your Switch" on page 66 for more information.
3. On the Mist portal, click the **Switches** tab on the left to navigate to the Switches page.
4. Click the Virtual Chassis to which you want to add the new member switch.
   The switch details page appears.
5. On the switch details page, click **Modify Virtual Chassis**.
   The Modify Virtual Chassis window appears.
6. On the Modify Virtual Chassis window, click **Add Switch**.

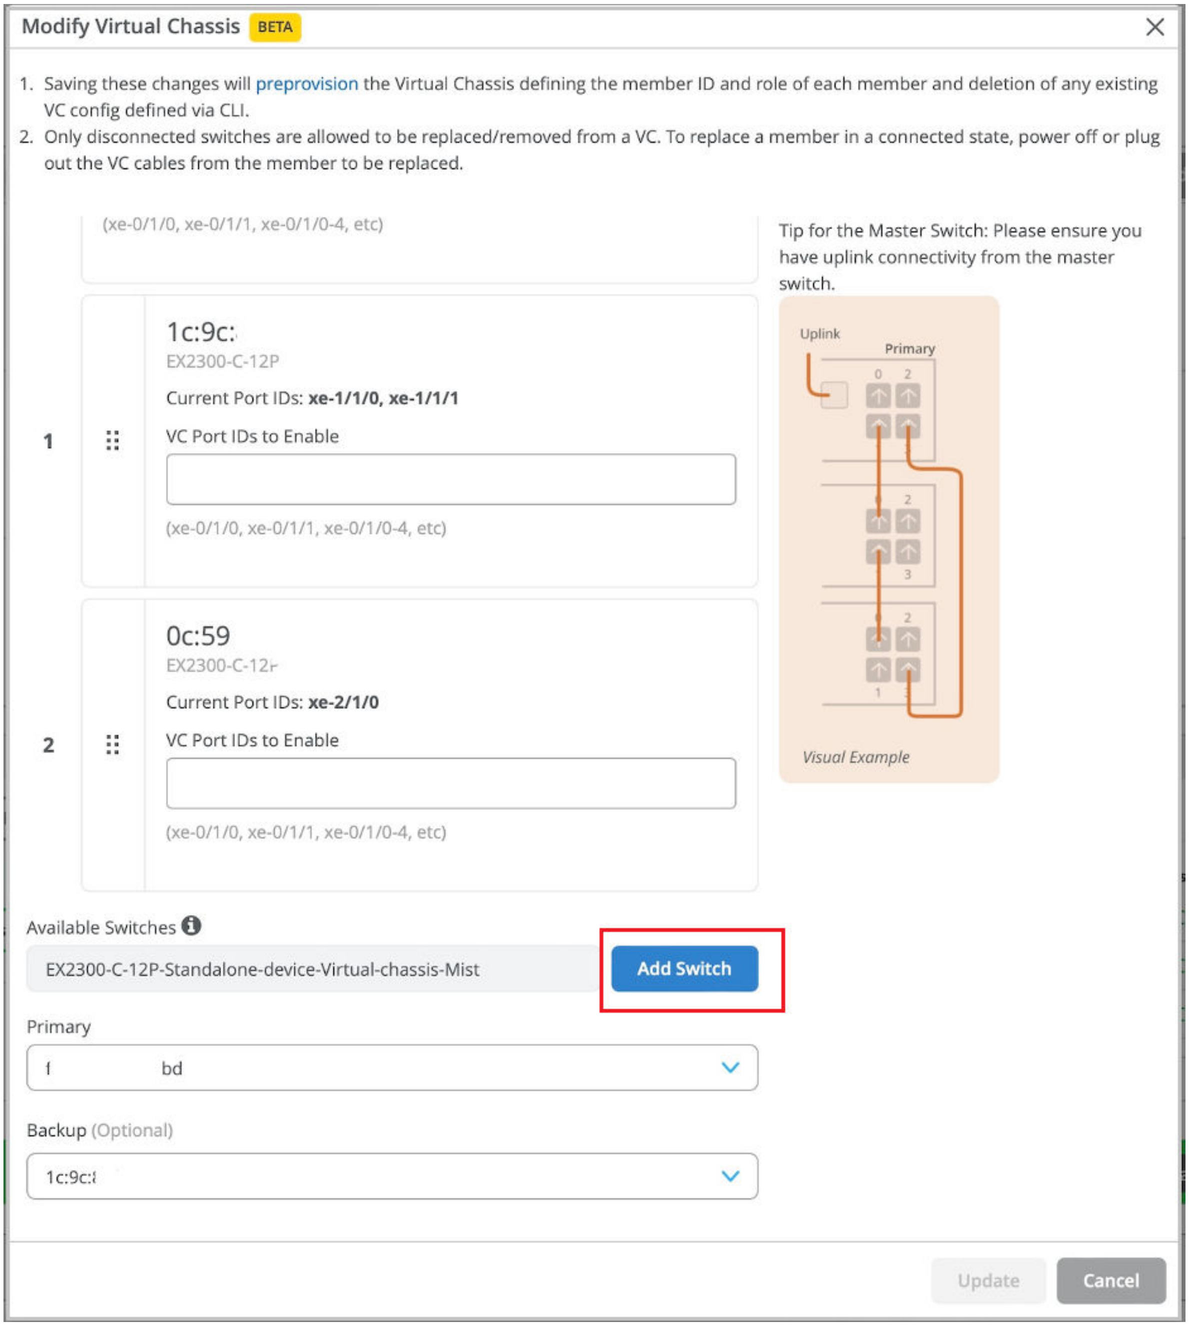

Modify Virtual Chassis BETA

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.

(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

Tip for the Master Switch: Please ensure you have uplink connectivity from the master switch.

1 — 1c:9c:
EX2300-C-12P
Current Port IDs: **xe-1/1/0, xe-1/1/1**
VC Port IDs to Enable
(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

2 — 0c:59
EX2300-C-12P
Current Port IDs: **xe-2/1/0**
VC Port IDs to Enable
(xe-0/1/0, xe-0/1/1, xe-0/1/0-4, etc)

Available Switches: EX2300-C-12P-Standalone-device-Virtual-chassis-Mist

Primary: f bd

Backup (Optional): 1c:9c:

7. Specify the VC port ID for the switch, if needed (the port ID configuration applies to the EX2300, EX4650, and QFX5120 switches).
8. Click **Update**.
9. 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'.

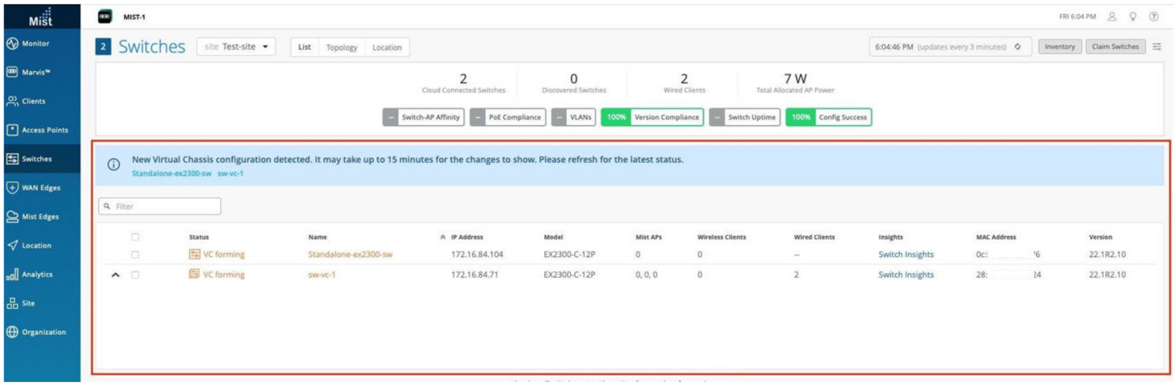

New Virtual Chassis configuration detected. It may take up to 15 minutes for the changes to show. Please refresh for the latest status.
Standalone-ex2300-sw sw-vc-1

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.

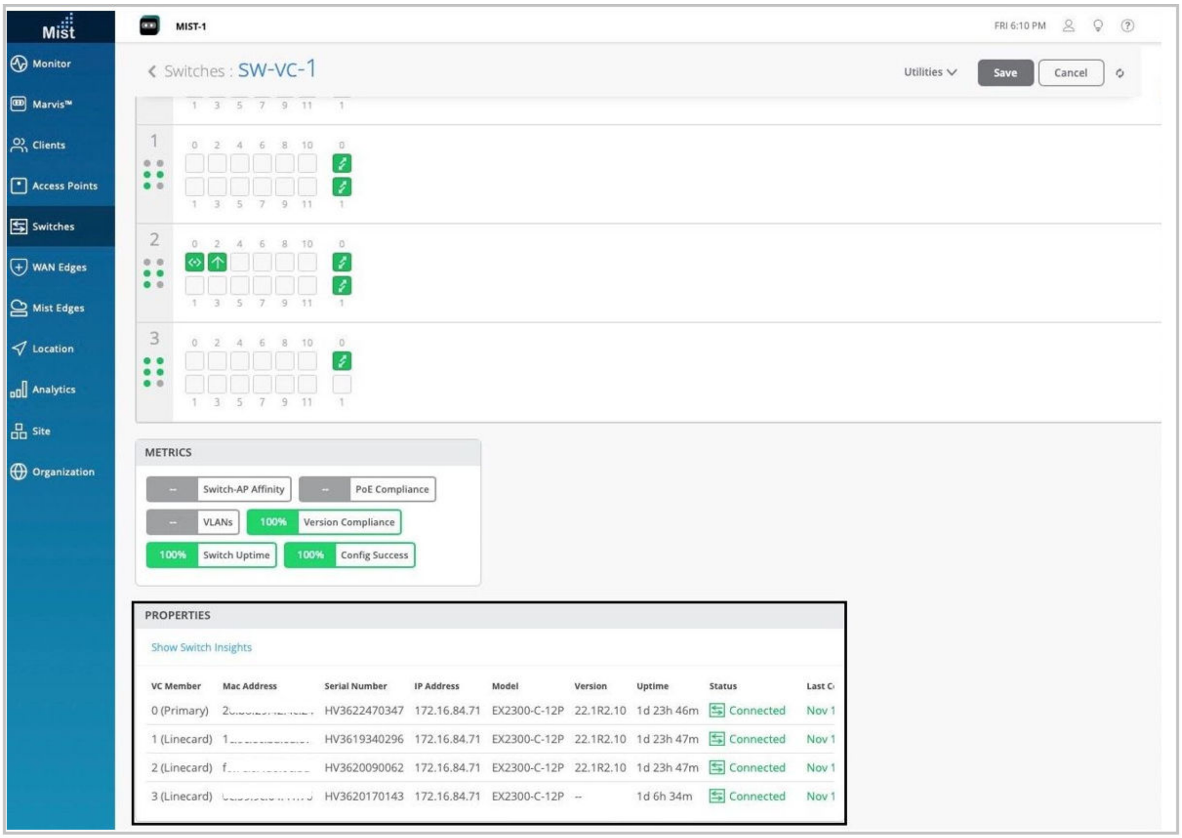

PROPERTIES

Show Switch Insights

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last C
0 (Primary)	2[redacted]	HV3622470347	172.16.84.71	EX2300-C-12P	22.1R2.10	1d 23h 46m	Connected	Nov 1
1 (Linecard)	1[redacted]	HV3619340296	172.16.84.71	EX2300-C-12P	22.1R2.10	1d 23h 47m	Connected	Nov 1
2 (Linecard)	f[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

5

CHAPTER

Campus Fabric Configuration

Which Campus Fabric Topology to Choose

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.

For an overview of campus fabric, watch the following video:

Video: Campus Fabric Overview

You can build and manage a campus fabric using the Mist portal. This topic describes the following campus fabric topologies, all of which Juniper Mist™ supports.

- EVPN Multihoming
- Campus Fabric Core-Distribution
- Campus Fabric IP Clos

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 4: Collapsed Core Topology

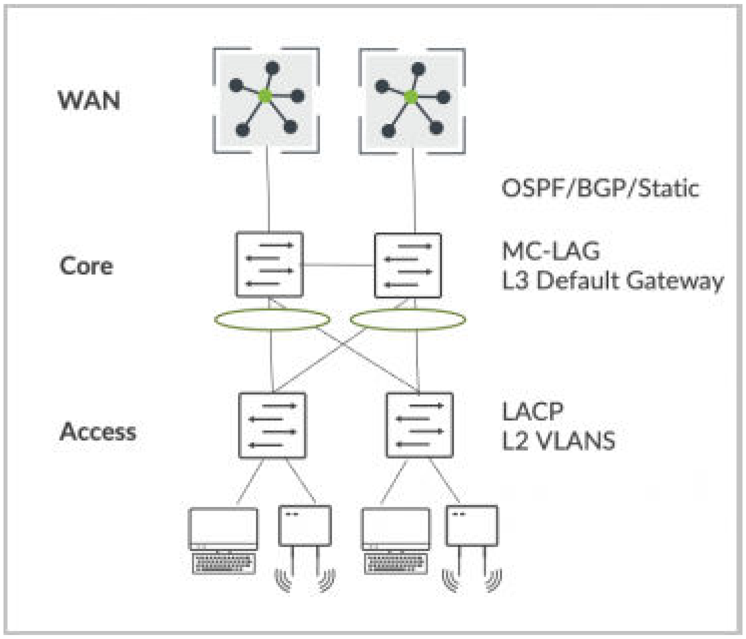

However, the traditional collapsed core deployment model presents the following challenges:

- 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 5: EVPN Multihoming

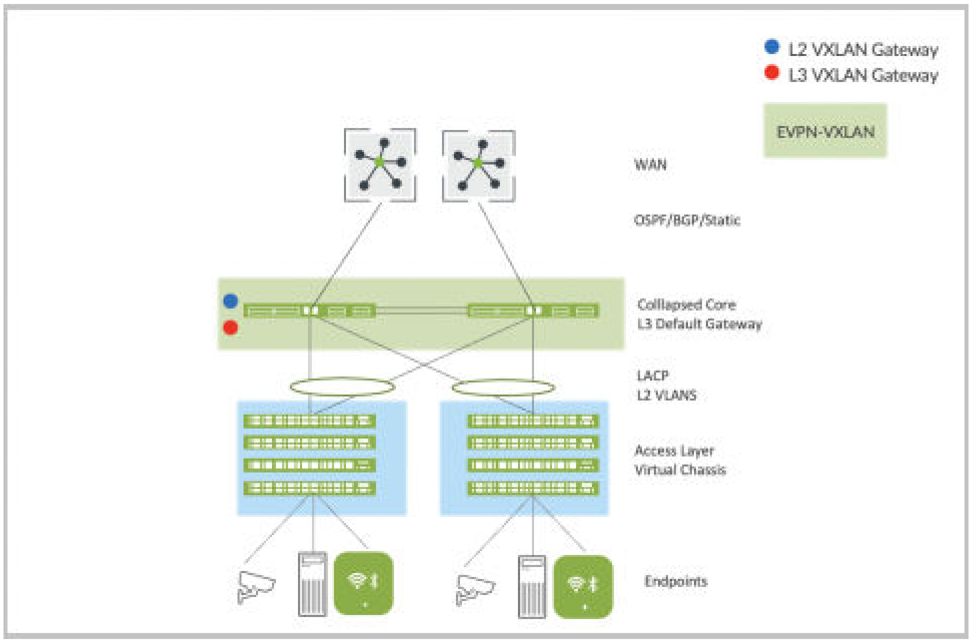

Choose EVPN Multihoming if you want to:

- 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: EX9200, EX4400-48F, and EX4400-24X, EX4650, QFX5120, QFX5110, QFX5700, and QFX5130

- 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 6: 3-Stage Core-Distribution-Access Network

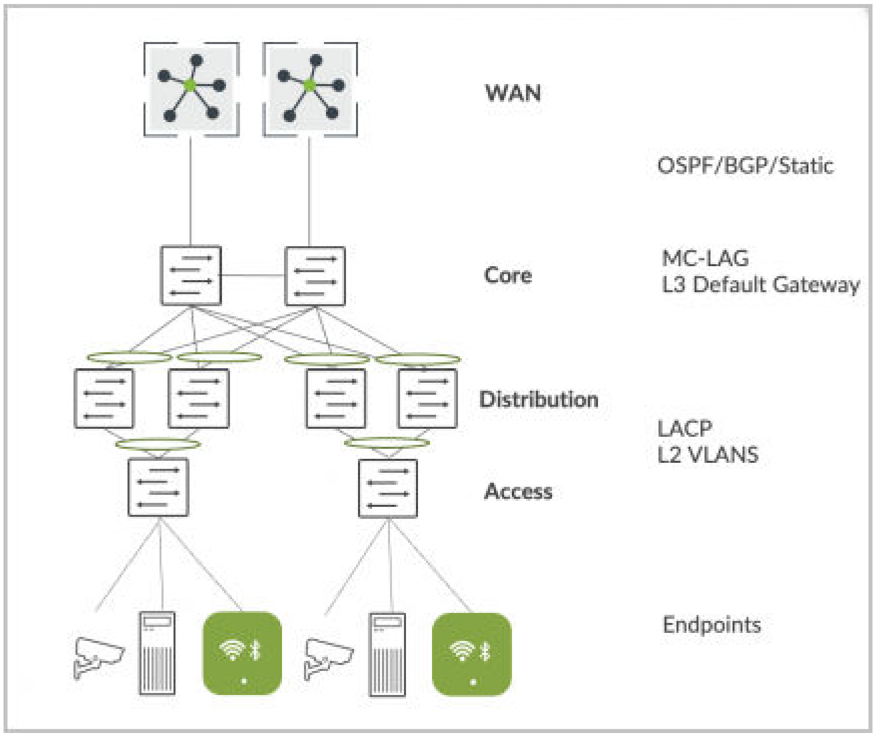

This traditional deployment model faces the following challenges:

- 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.

Figure 7: Campus Fabric Core-Distribution - CRB or ERB

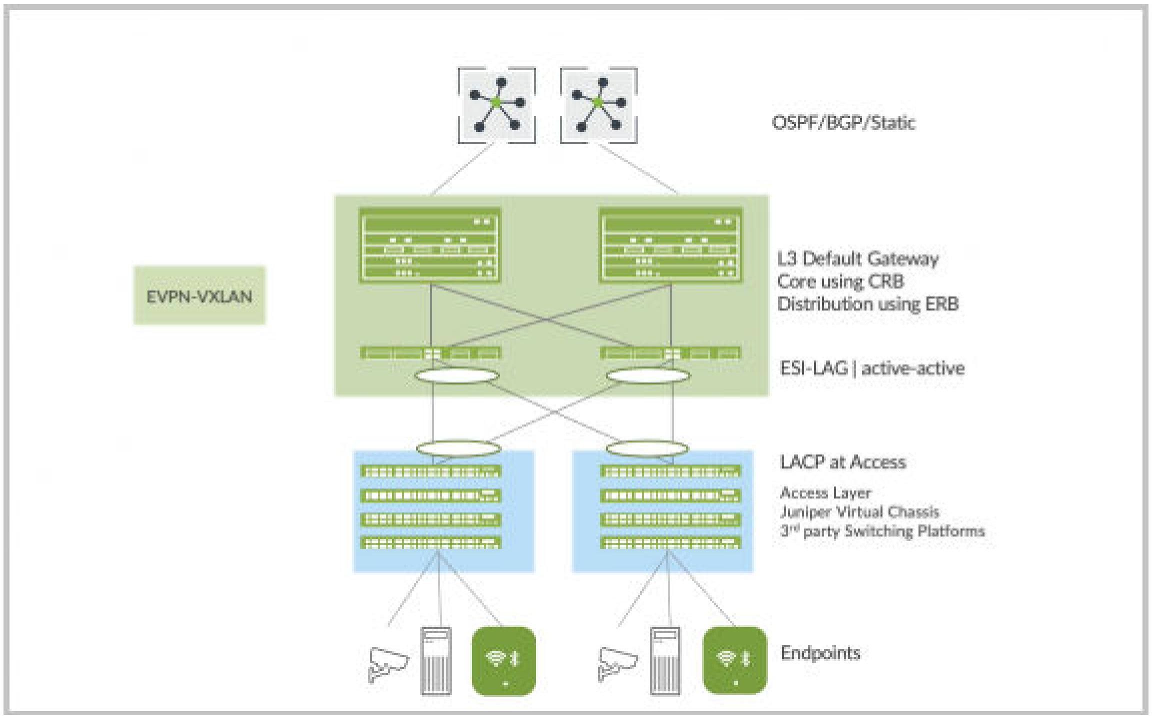

Choose Campus Fabric Core-Distribution if you want to:

- 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
- 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

Figure 8: Campus Fabric IP Clos 5 Stage

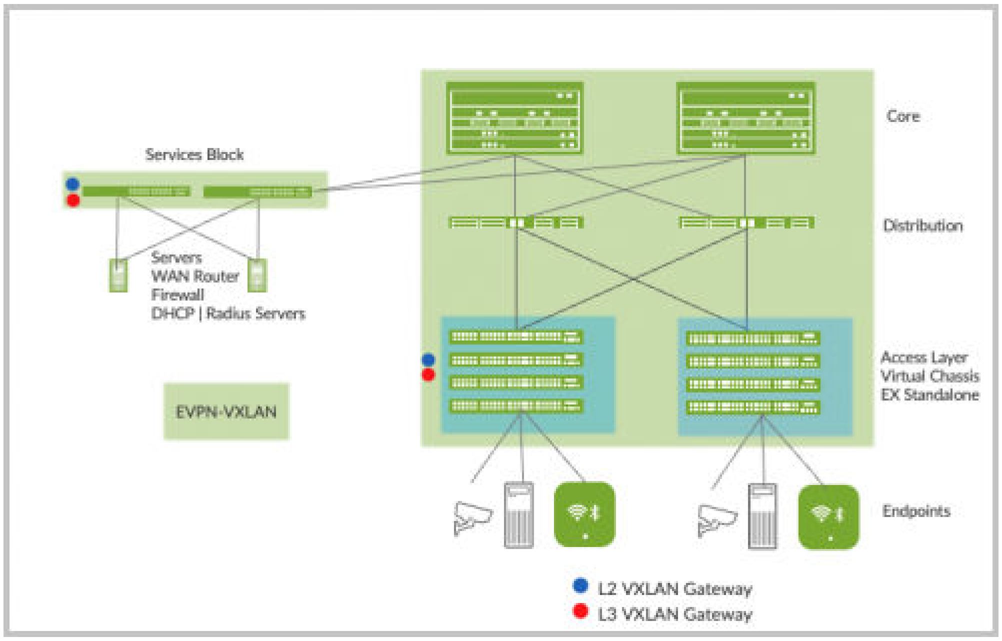

Figure 9: Campus Fabric IP Clos 3 Stage

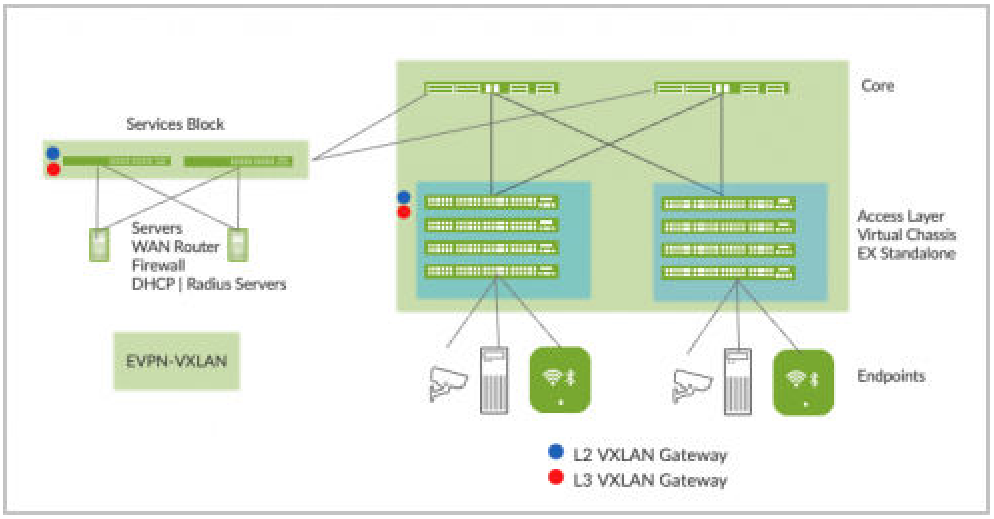

The following Juniper Network platforms support Campus Fabric IP Clos:

- 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, QFX5170, EX9200, and QFX10k

Configure Campus Fabric 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.

For a detailed configuration example, see Campus Fabric EVPN Multihoming Workflow.

Video: Video: Deployment of Campus Fabric EVPN Multihoming With Wired Assurance

Video:

To configure campus fabric EVPN multihoming:

1. Click **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.

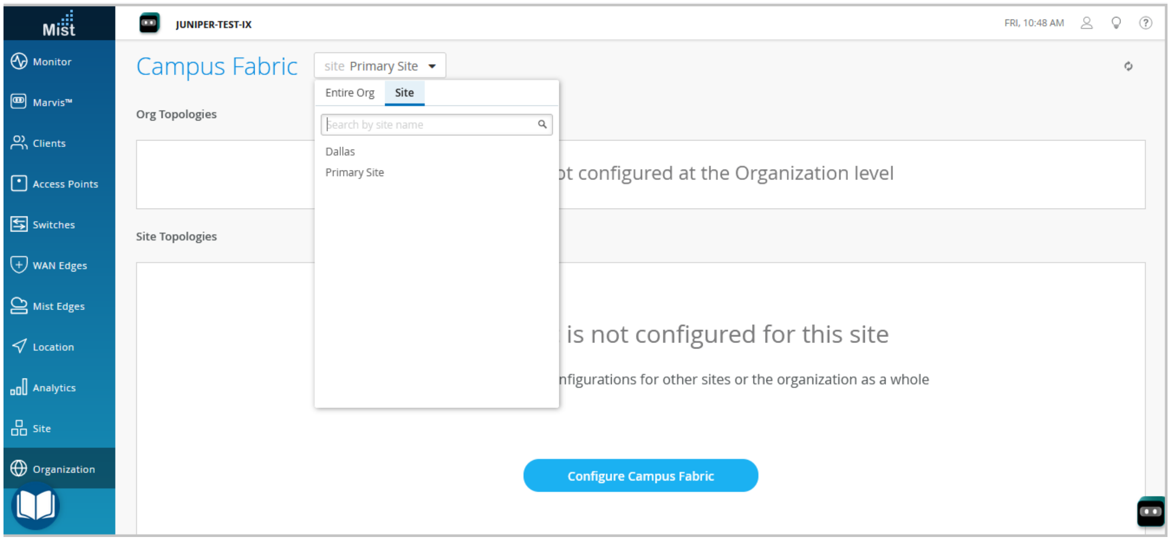

ot configured at the Organization level

is not configured for this site

nfigurations for other sites or the organization as a whole

The topology type EVPN Multihoming is available only for the site-specific campus fabric. You cannot build it for an entire organization.

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 **EVPN Multihoming**.

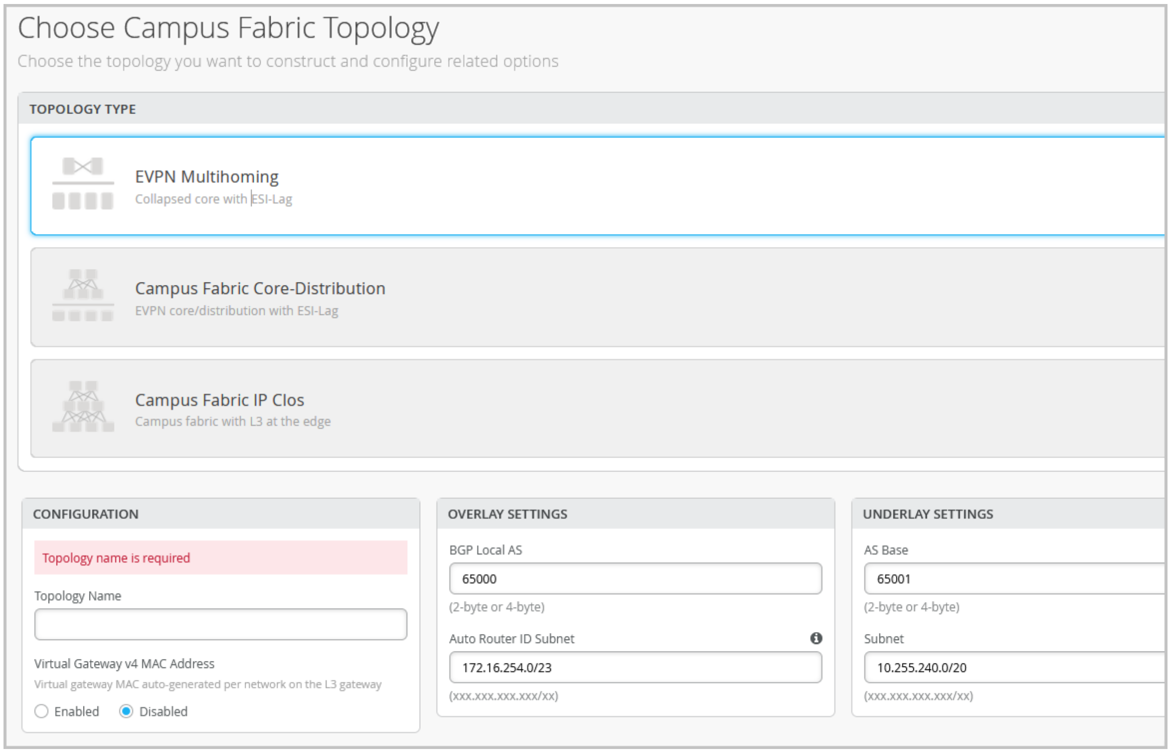

Choose Campus Fabric Topology

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)

Auto Router ID Subnet
172.16.254.0/23
(xxx.xxx.xxx.xxx/xx)

UNDERLAY SETTINGS

AS Base
65001
(2-byte or 4-byte)

Subnet
10.255.240.0/20
(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**—If you enable it, Mist provides a unique MAC address to each Layer 3 (L3) virtual gateway (per network). This setting is 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 to ensure that the AS numbers are never advertised outside the fabric.

- **Auto Router ID Subnet**—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.

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.

- **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.

6. Click **Continue** to go to the **Nodes** tab, where you can select devices that form part of the campus fabric deployment.

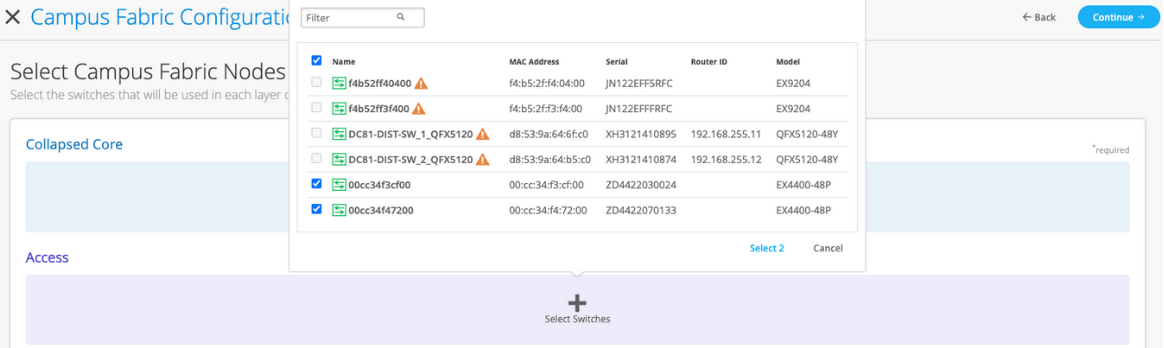

Campus Fabric Configurati

Select Campus Fabric Nodes

Select the switches that will be used in each layer c

☑	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.

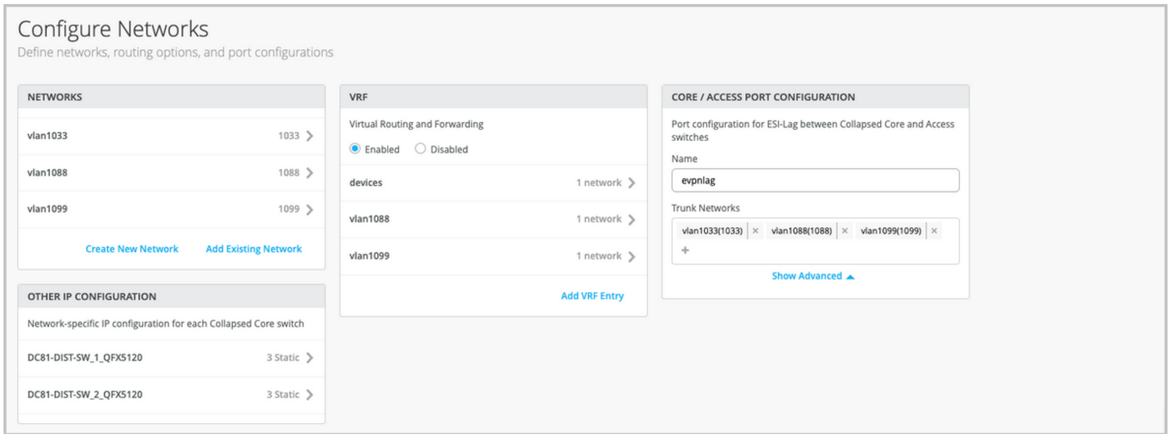

Configure Networks

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. 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 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.

      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 on the VRF tile. 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. 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.

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.

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.
- **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.

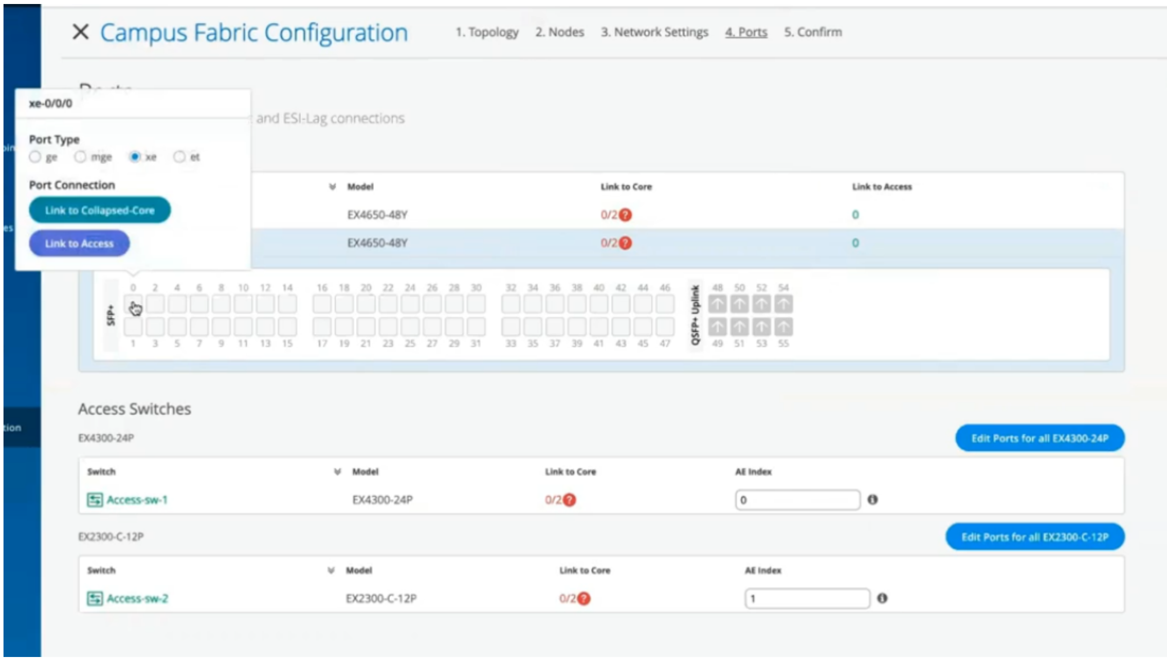

× Campus Fabric Configuration

1. Topology 2. Nodes 3. Network Settings 4. Ports 5. Confirm

and ESI-Lag connections

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.

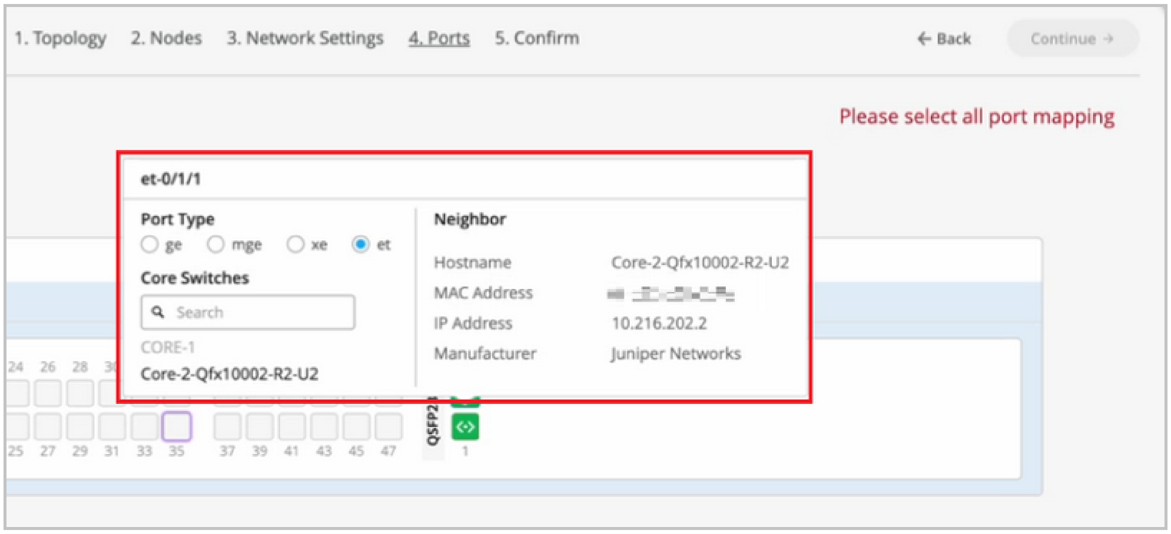

Please select all port mapping

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).

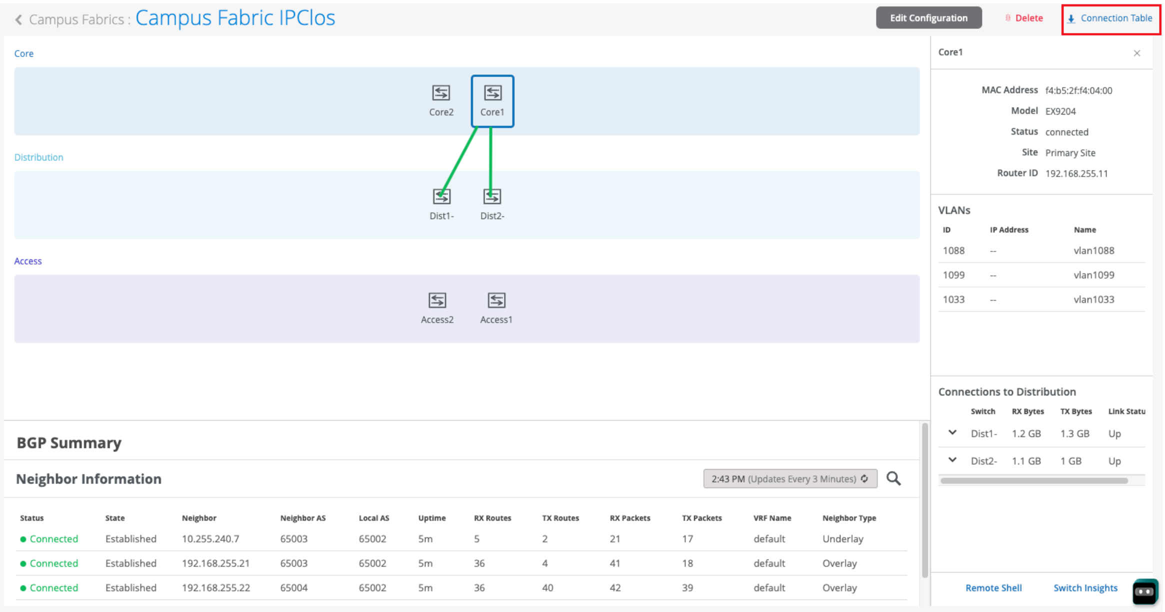

Core1

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.

Configure Campus Fabric Core-Distribution

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)

Video: Video: Deployment of Campus Fabric Core-Distribution

Video:

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.

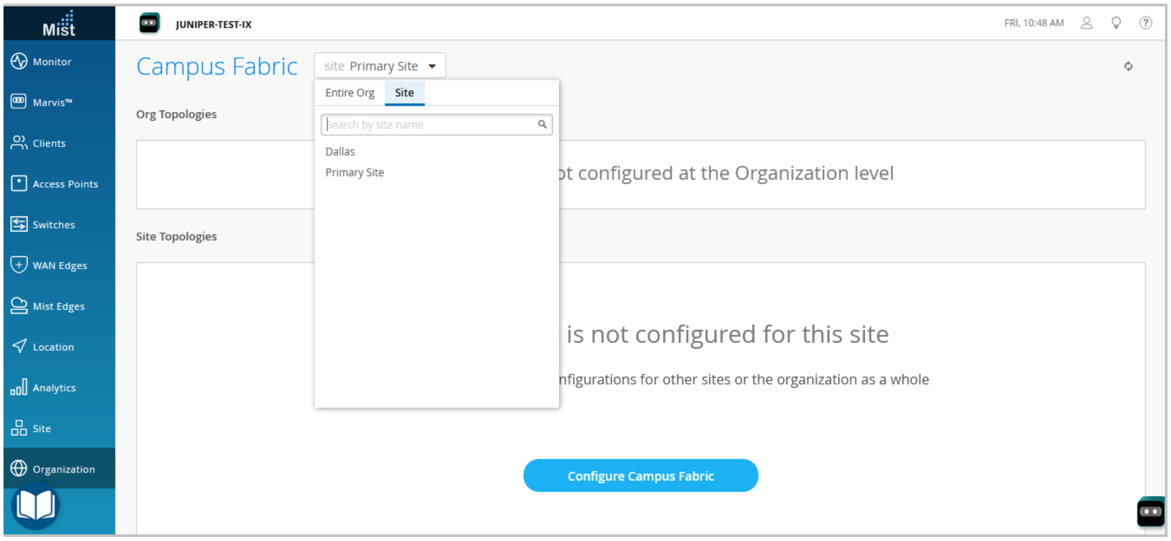

ot configured at the Organization level

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**.

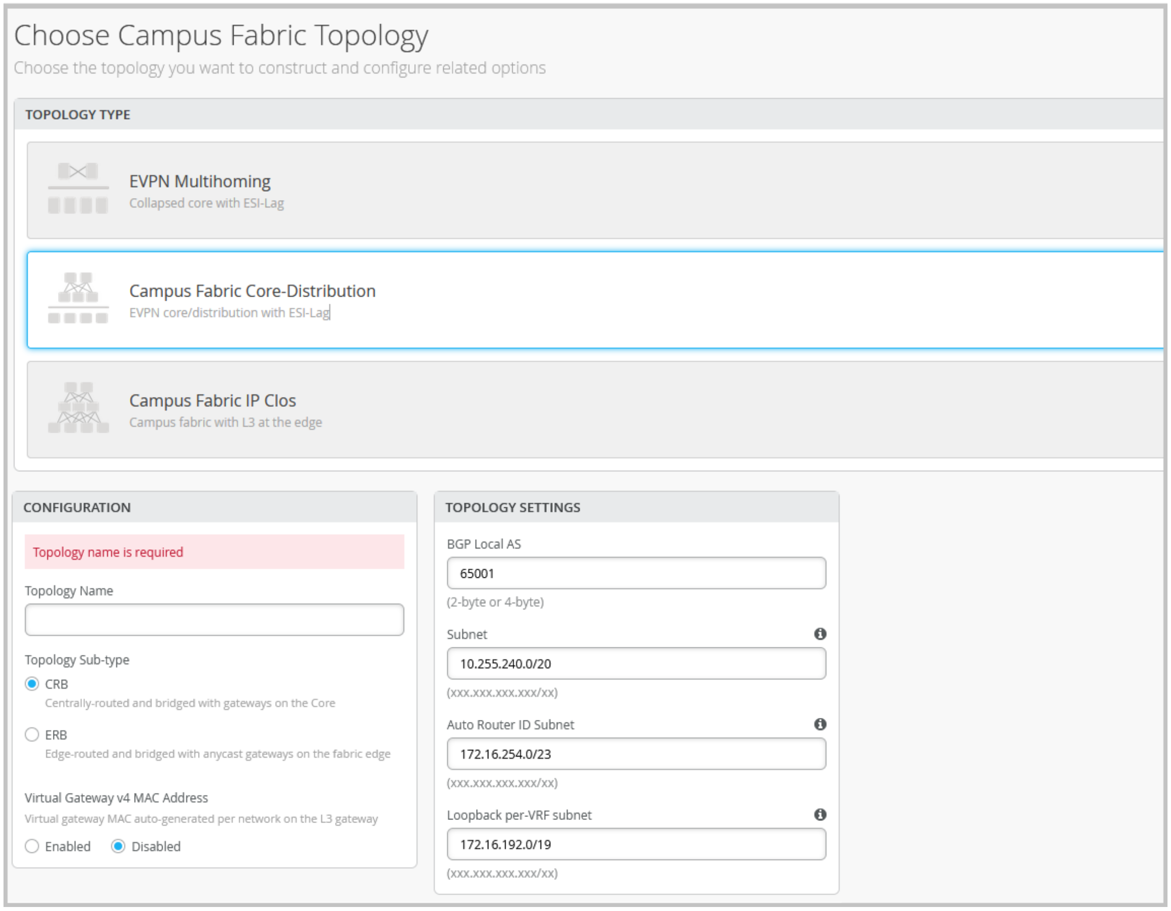

Choose Campus Fabric Topology

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

Virtual Gateway v4 MAC Address
Virtual gateway MAC auto-generated per network on the L3 gateway

○ Enabled ● Disabled

TOPOLOGY SETTINGS

BGP Local AS: 65001
(2-byte or 4-byte)

Subnet: 10.255.240.0/20
(xxx.xxx.xxx.xxx/xx)

Auto Router ID Subnet: 172.16.254.0/23
(xxx.xxx.xxx.xxx/xx)

Loopback per-VRF subnet: 172.16.192.0/19
(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**—Available only if you have selected CRB. If you enable it, Mist provides a unique MAC address to each L3 IRB interface (per network).

    - **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 to ensure that the AS numbers are never advertised outside the fabric.

- **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.

- **Auto Router ID Subnet**—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 (172.16.192.0/19), which you can modify. 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.

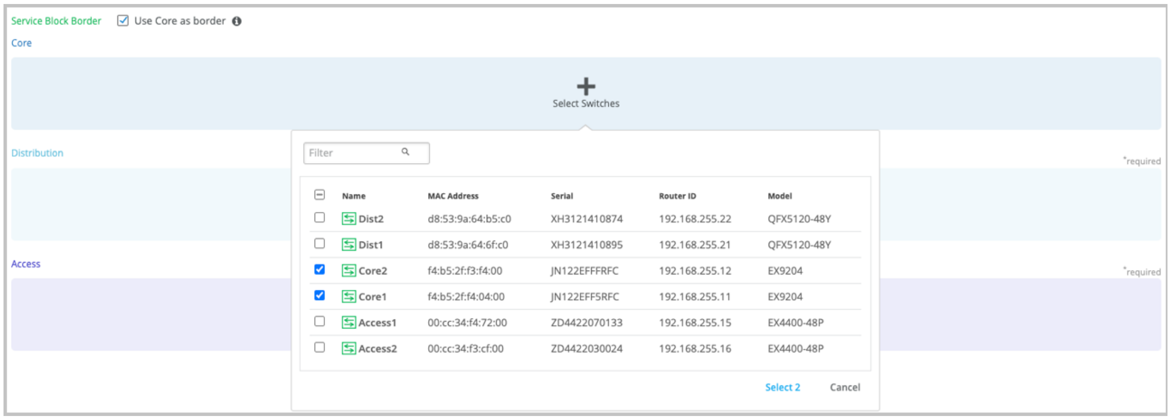

⊟	Name	MAC Address	Serial	Router ID	Model
□	Dist2	d8:53:9a:64:b5:c0	XH3121410874	192.168.255.22	QFX5120-48Y
□	Dist1	d8:53:9a:64:6f:c0	XH3121410895	192.168.255.21	QFX5120-48Y
☑	Core2	f4:b5:2f:f3:f4:00	JN122EFFFRFC	192.168.255.12	EX9204
☑	Core1	f4:b5:2f:f4:04:00	JN122EFF5RFC	192.168.255.11	EX9204
□	Access1	00:cc:34:f4:72:00	ZD4422070133	192.168.255.15	EX4400-48P
□	Access2	00:cc:34:f3:cf:00	ZD4422030024	192.168.255.16	EX4400-48P

For organization-level topologies, you can create pods to group access and distribution devices. 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.

7. Add switches to the Core, Distribution, and Access layer sections.

   To add switches:

   a. Click **Select 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 service 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.

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:

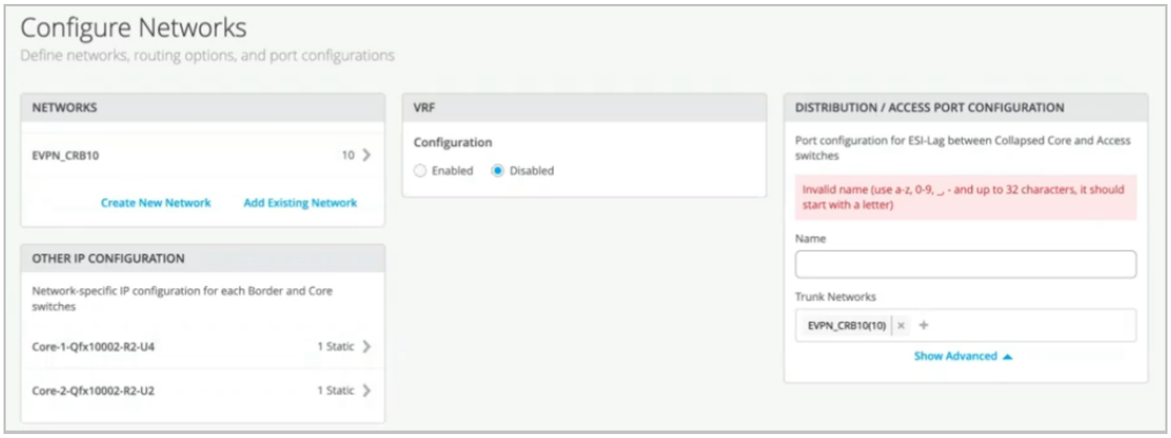

Configure Networks

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.

   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 on the VRF tile. 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. 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.

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.

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.

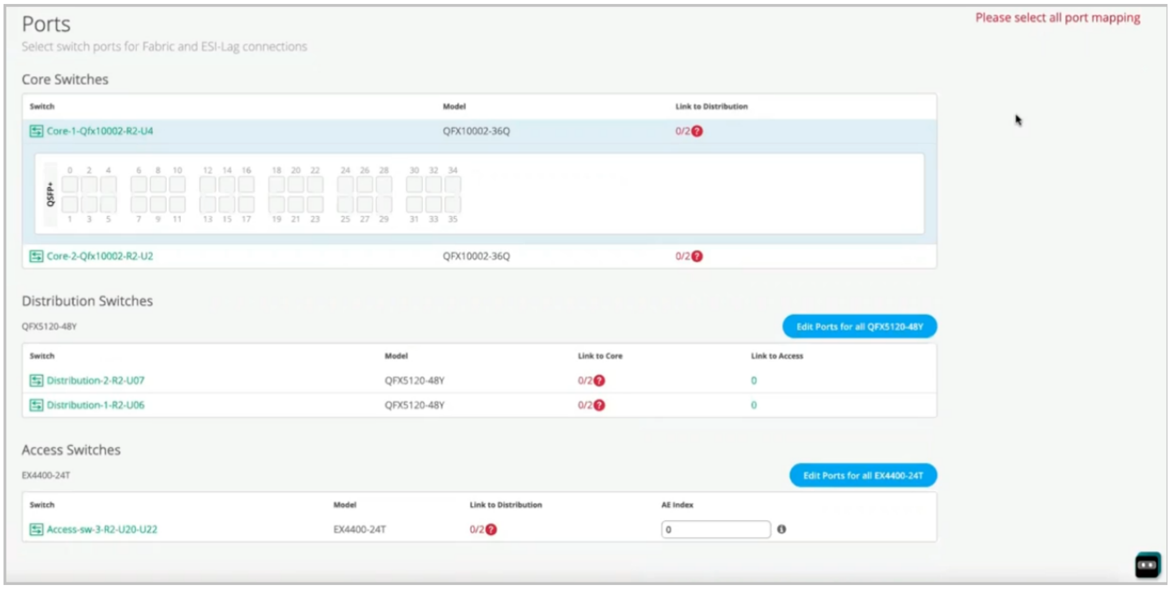

Ports

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.

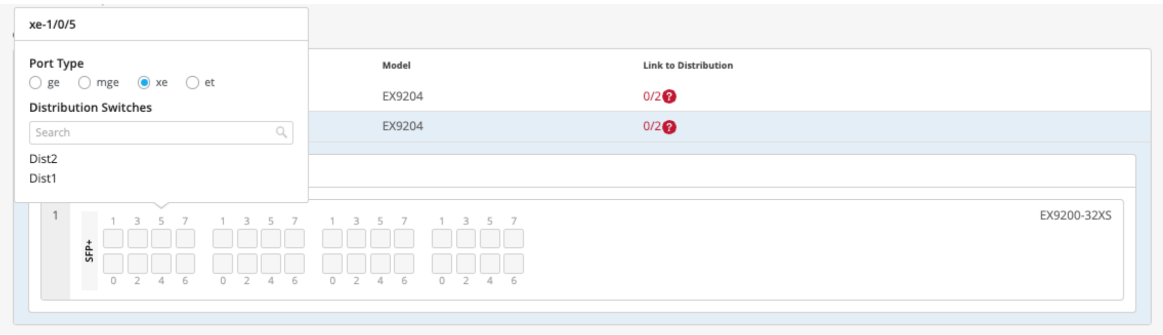

xe-1/0/5

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.

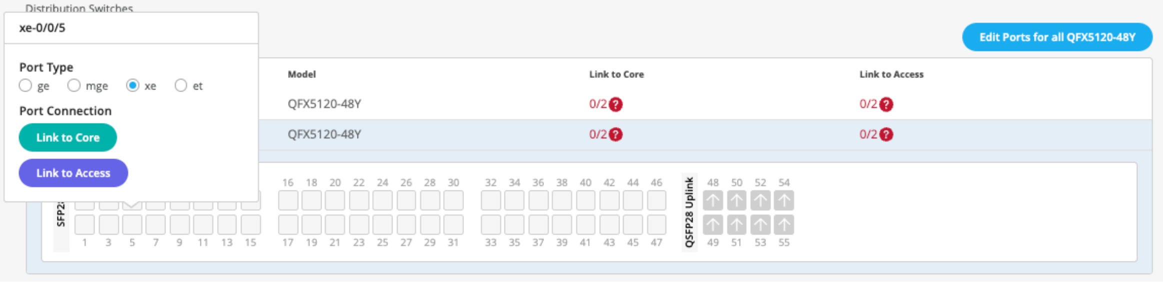

Distribution Switches

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.

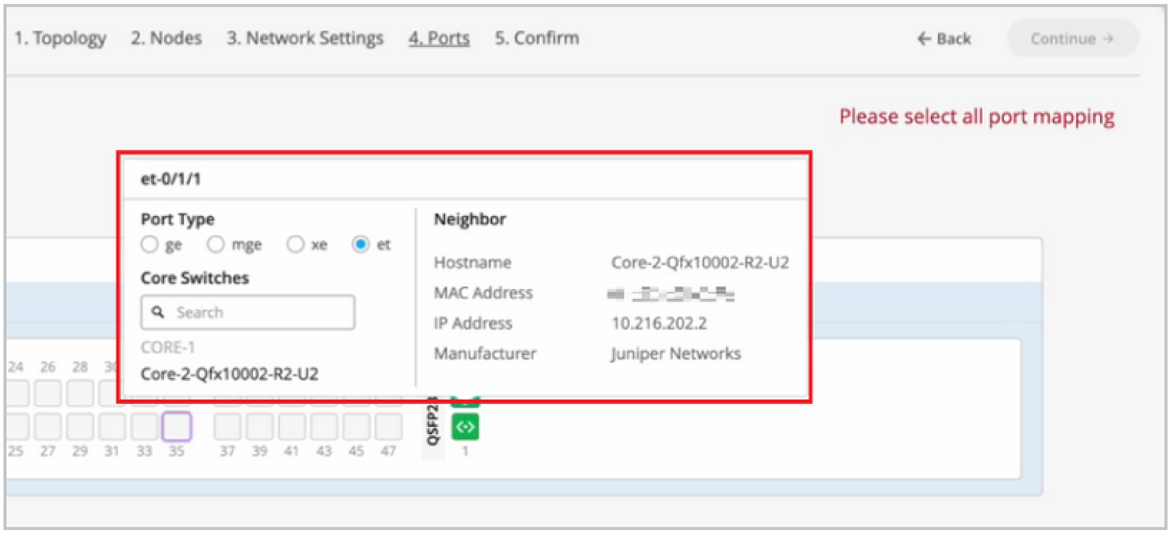

Please select all port mapping

et-0/1/1

Port Type
○ ge ○ mge ○ xe ● et

Core Switches

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).

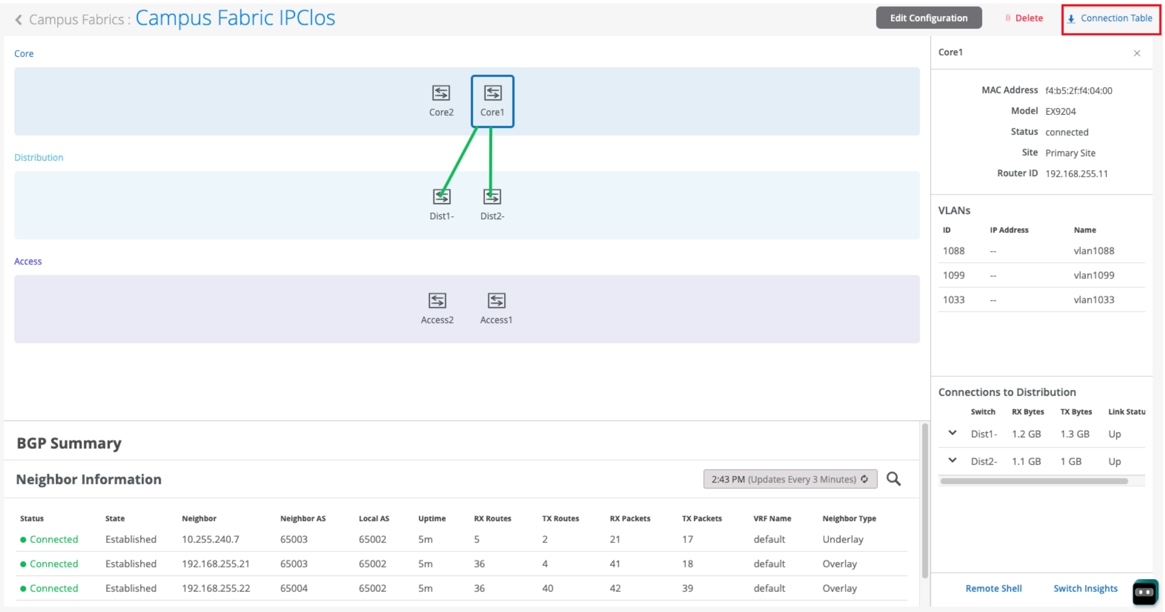

Core1

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 of Campus Fabric Core Distribution CRB (JVD) and Campus Fabric Core-Distribution ERB (JVD).

Configure Campus Fabric IP Clos

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.

Video: Video: Deployment of Campus Fabric IPCLOS

Video:

For more detailed information about IP Clos architecture and its deployment, see Campus Fabric IP Clos Wired Assurance.

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.

   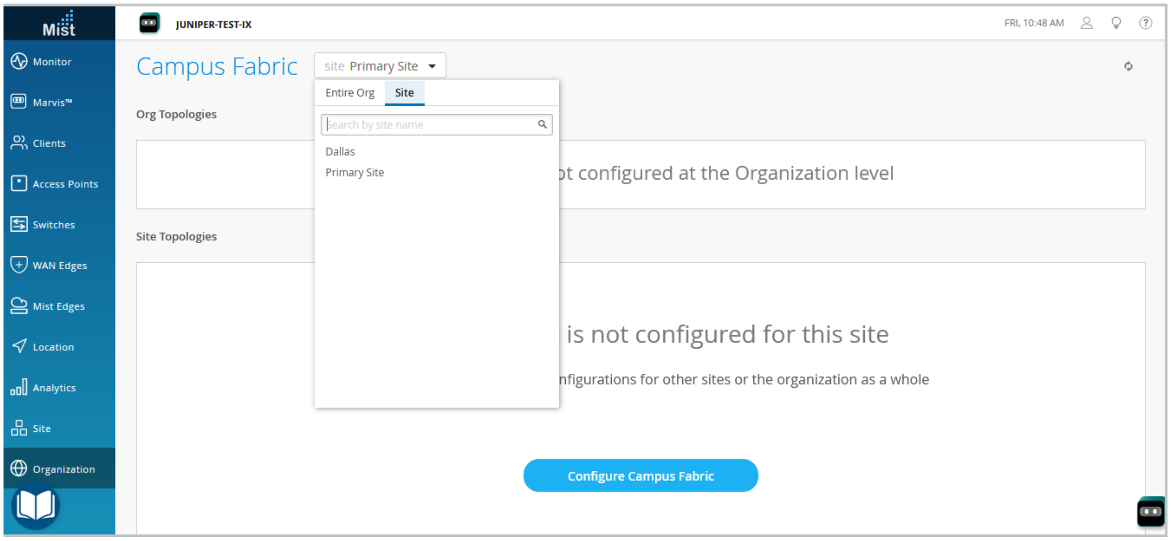

   ot configured at the Organization level

   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**.

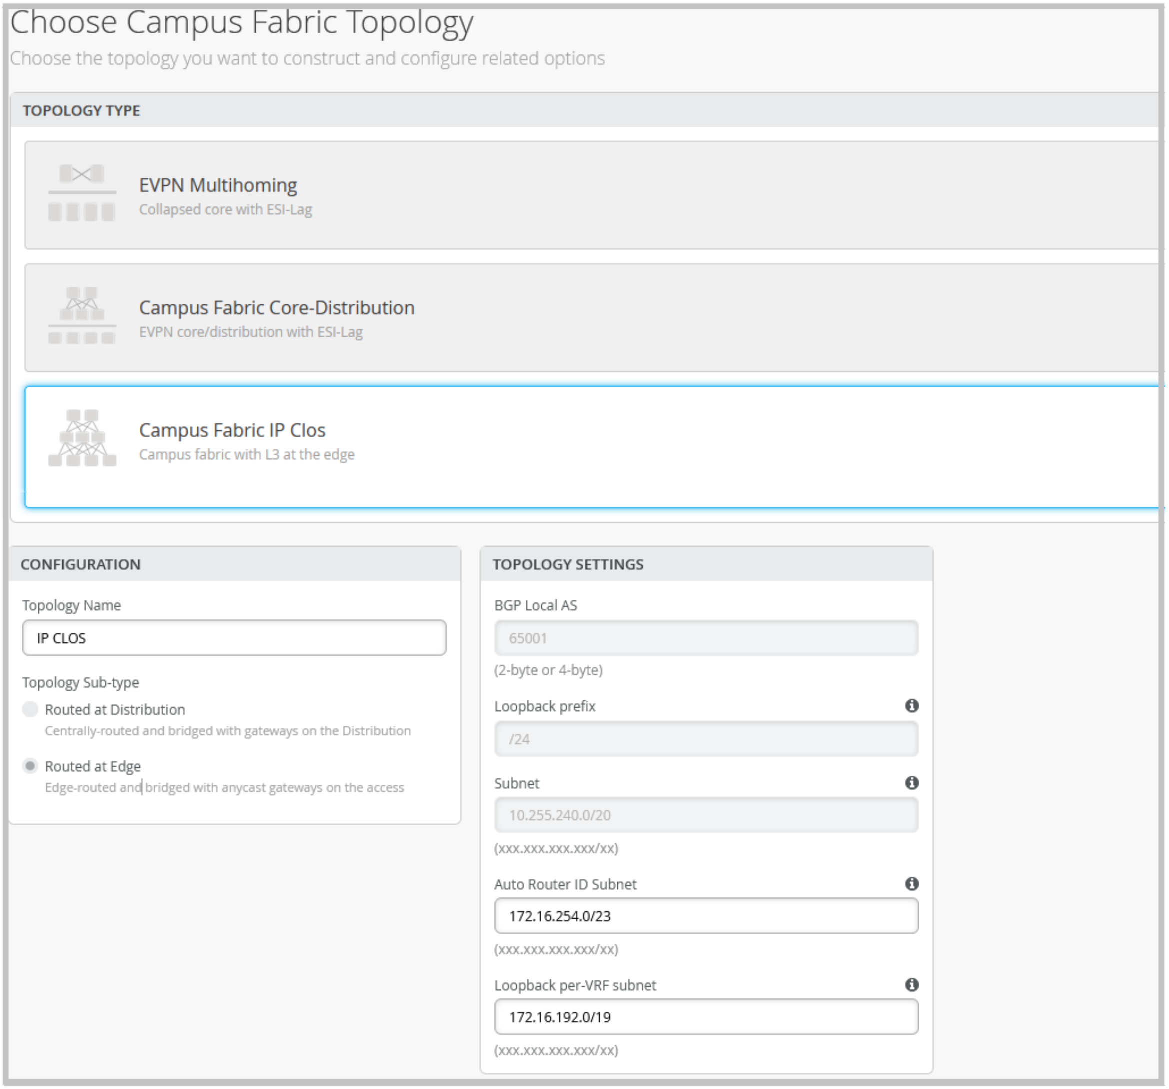

Choose Campus Fabric Topology

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 Name: IP CLOS

Topology Sub-type

○ Routed at Distribution
Centrally-routed and bridged with gateways on the Distribution

● Routed at Edge
Edge-routed and bridged with anycast gateways on the access

TOPOLOGY SETTINGS

BGP Local AS: 65001
(2-byte or 4-byte)

Loopback prefix: /24

Subnet: 10.255.240.0/20
(xxx.xxx.xxx.xxx/xx)

Auto Router ID Subnet: 172.16.254.0/23
(xxx.xxx.xxx.xxx/xx)

Loopback per-VRF subnet: 172.16.192.0/19
(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:
  - **Routed at Distribution**—If you select this option, Mist provisions layer 3 (L3) integrated routing and bridging (IRB) interfaces on the distribution layer. The distribution switches are configured with the same IP address for each L3 subnet. This option aligns with north-

south traffic patterns and configures the access layer as an L2 VXLAN gateway. This option is preferred for higher-scale deployments.

    - **Virtual Gateway v4 MAC Address**—Available only if you have selected Routed at Distribution. If you enable it, Mist provides a unique MAC address to each L3 IRB interface (per network).

  - **Routed at Edge**—(Default) 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. This option utilizes anycast addressing for all devices participating in the L3 subnet. This option provides a smaller blast radius for broadcast traffic and is ideal for east-west traffic patterns and IP Multicast environments.

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 to ensure that the AS numbers are never advertised outside the fabric.

  - **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.

  - **Auto Router ID Subnet**—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 (172.16.192.0/19), which you can modify. 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.

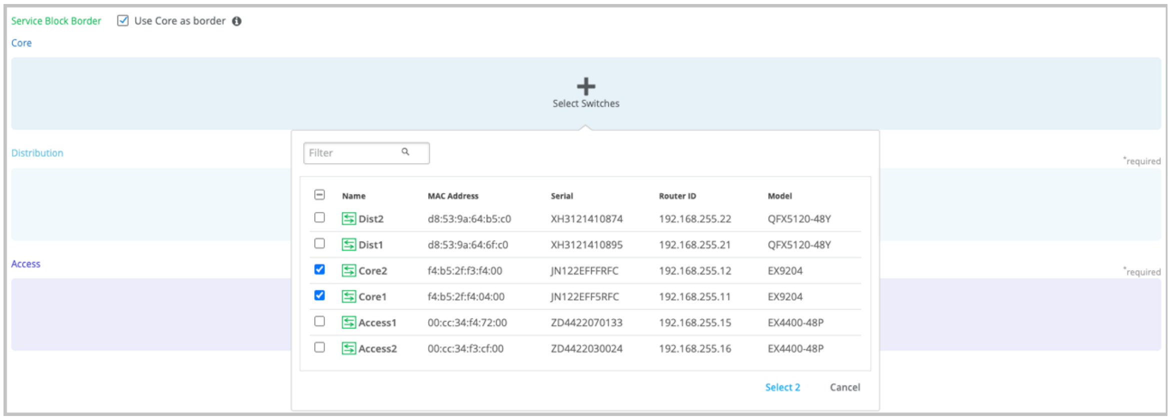

⊟	Name	MAC Address	Serial	Router ID	Model
☐	Dist2	d8:53:9a:64:b5:c0	XH3121410874	192.168.255.22	QFX5120-48Y
☐	Dist1	d8:53:9a:64:6f:c0	XH3121410895	192.168.255.21	QFX5120-48Y
☑	Core2	f4:b5:2f:f3:f4:00	JN122EFFFRFC	192.168.255.12	EX9204
☑	Core1	f4:b5:2f:f4:04:00	JN122EFF5RFC	192.168.255.11	EX9204
☐	Access1	00:cc:34:f4:72:00	ZD4422070133	192.168.255.15	EX4400-48P
☐	Access2	00:cc:34:f3:cf:00	ZD4422030024	192.168.255.16	EX4400-48P

7. Add switches to the Core, Distribution, and Access layer sections.

   To add the switches:

   a. Click **Select 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 service 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.

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.

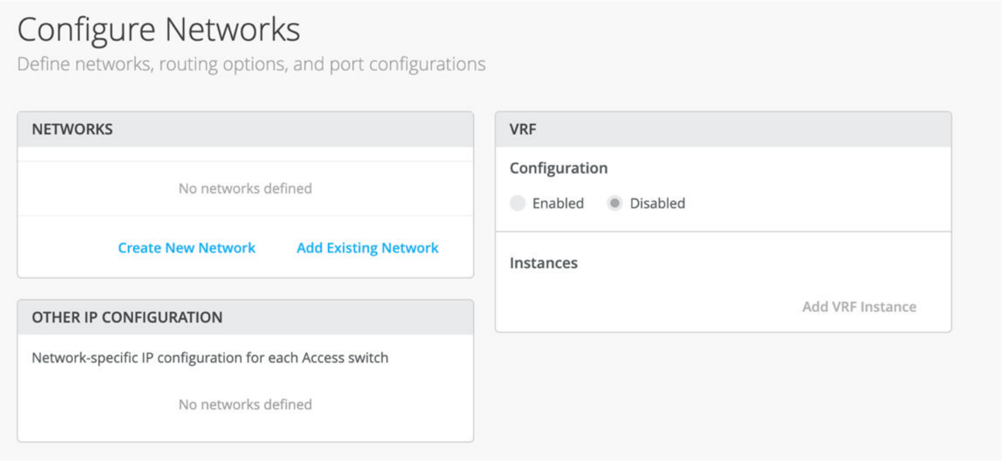

Configure Networks

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.

   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 in the VRF tile. 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. 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.

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.

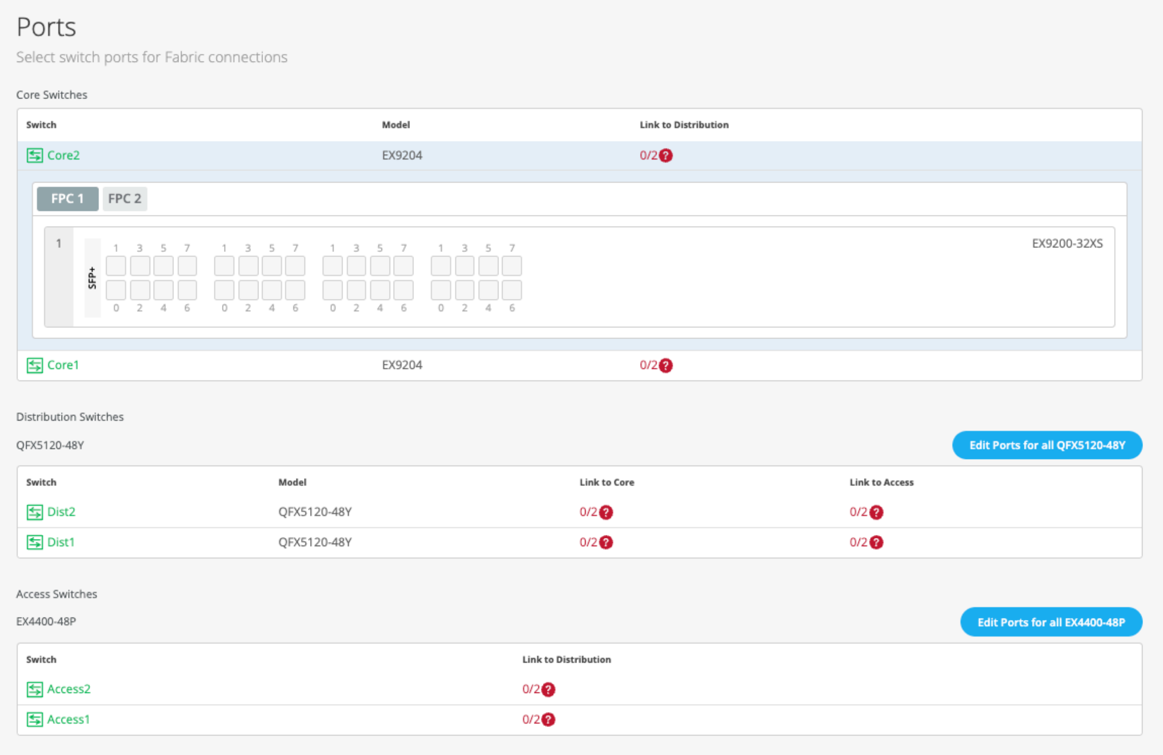

Ports

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.

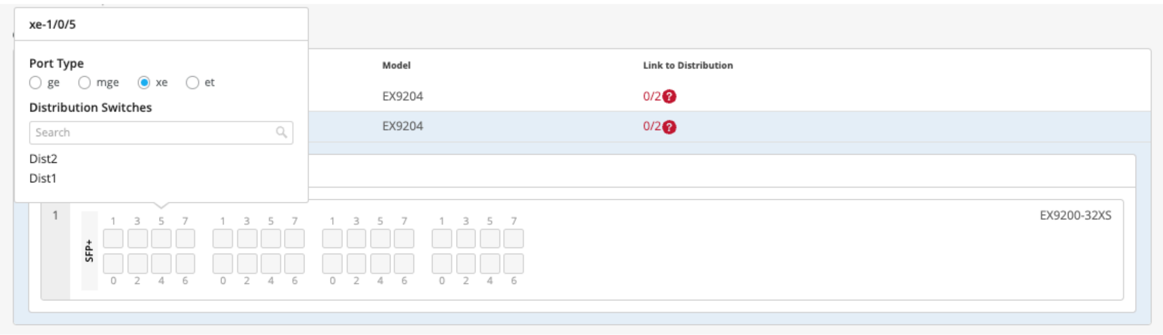

xe-1/0/5

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.

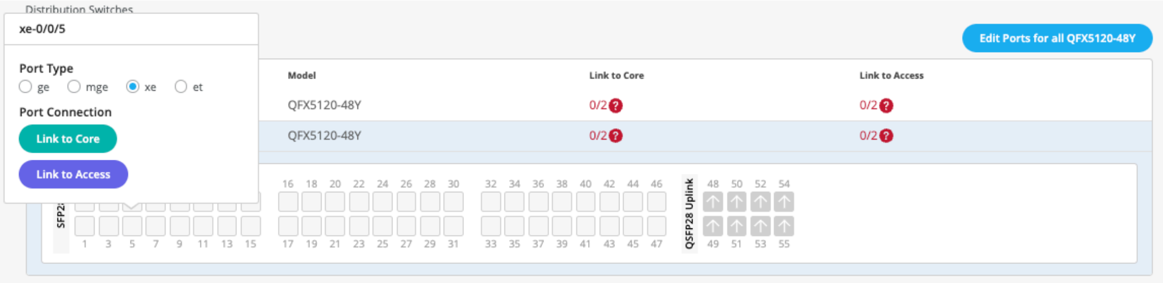

Distribution Switches

xe-0/0/5

Port Type

○ ge ○ mge ● xe ○ et

Port Connection

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 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.

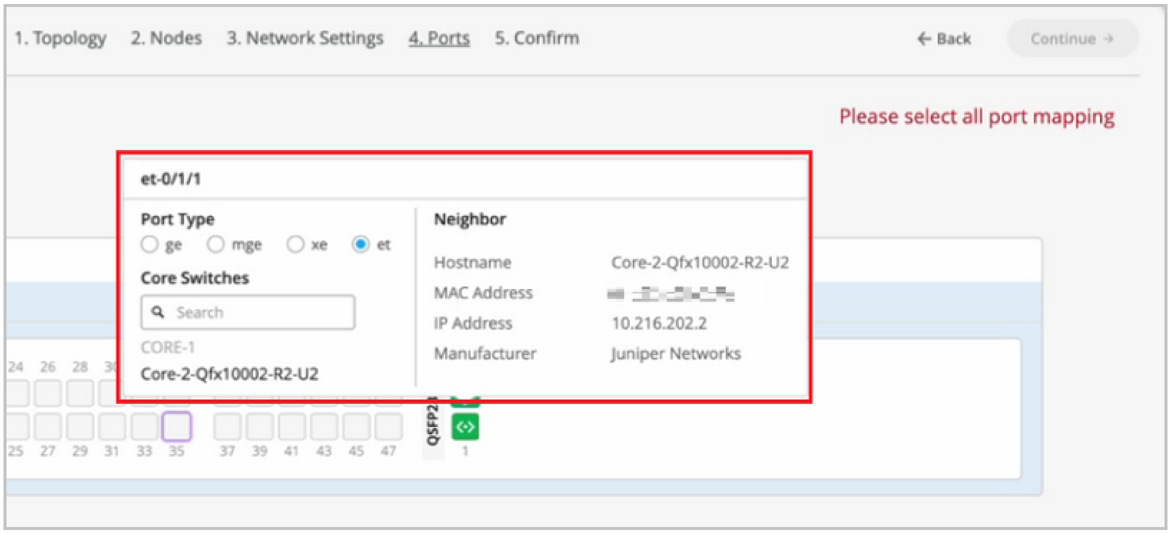

Please select all port mapping

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).

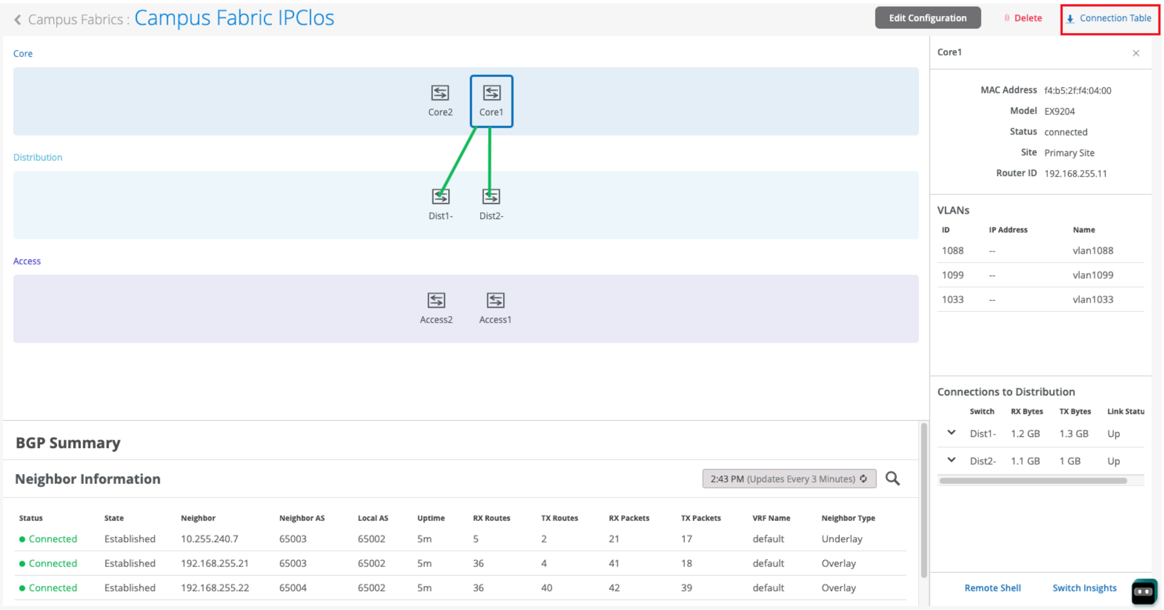

Core1

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.

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:

6

CHAPTER

Wired Service Levels

Wired Service-Level Expectations (SLEs)

IN THIS SECTION

Juniper Mist^{TM} cloud continuously collects network telemetry data and uses machine learning to analyze the end-user experience. 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. 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.

For a quick overview of Juniper Mist wired SLEs, watch the following video:

Video: Wired SLE Overview

View SLE Metrics

The wired SLE dashboards display the percentage of time that the SLE metrics met the specified service-level expectation goal within a specific time range. These metrics are categorized into classifiers and sub-classifiers, which provide additional details to identify the specific causes of failure. With this information, you can easily identify and address the issues affecting the end-user experience.

Mist Wired SLEs provide the following metrics to help you assess the end-user experience on your networks:

- Throughput
- Switch Health
- Successful Connects

To view the SLE metrics on the Wired SLE dashboard, click **Monitor** > **Service Levels**, and then select the **Wired** tab.

Figure 10: SLE Dashboard

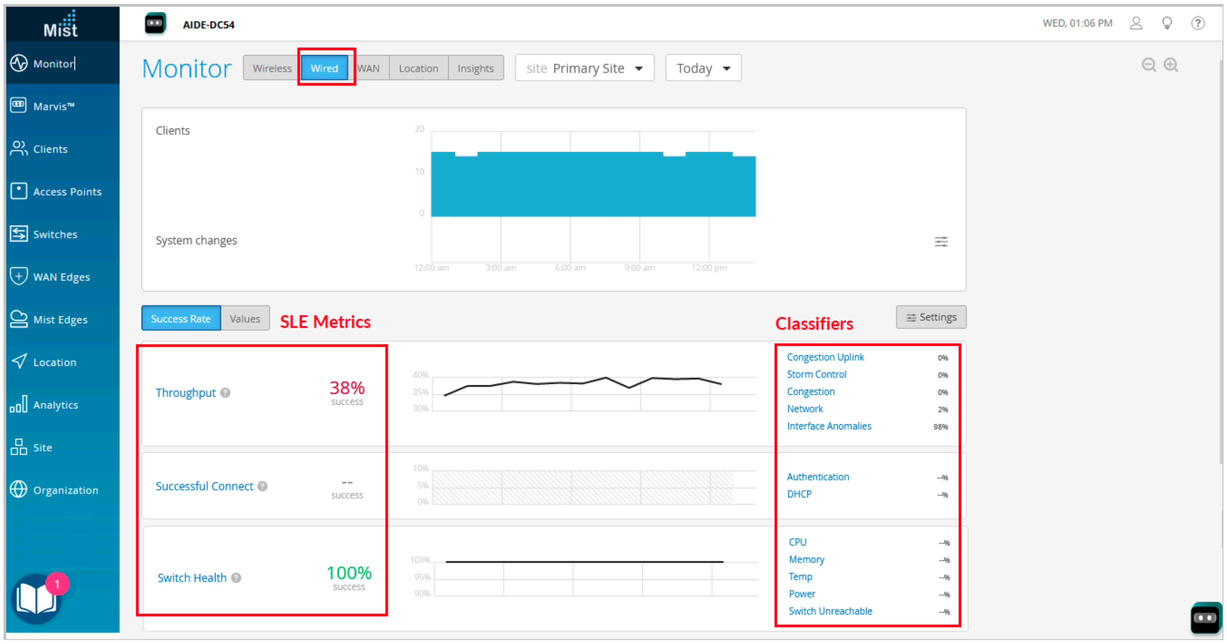

Each metric has classifiers and sub-classifiers that display information to help you identify failures and narrow down the specific problem area. To view the associated sub-classifiers, simply click a classifier. You'll see a tabbed view that includes:

- **Statistics**—Shows the overall success rate for the SLE metric.
- **Timeline**—Shows the timeline of the failures. For example, the dashboard can show the bad user minutes caused by issues belonging to a particular classifier over a period of time.
- **Distribution**—Shows the percentage of impact across different attributes such as interfaces, switches, VLANs, and clients.
- **Affected items—**Shows the specific items that failed to meet the service-level goal. Examples: switches, interfaces, and clients.

Here's an example of a **Throughput** metric view:

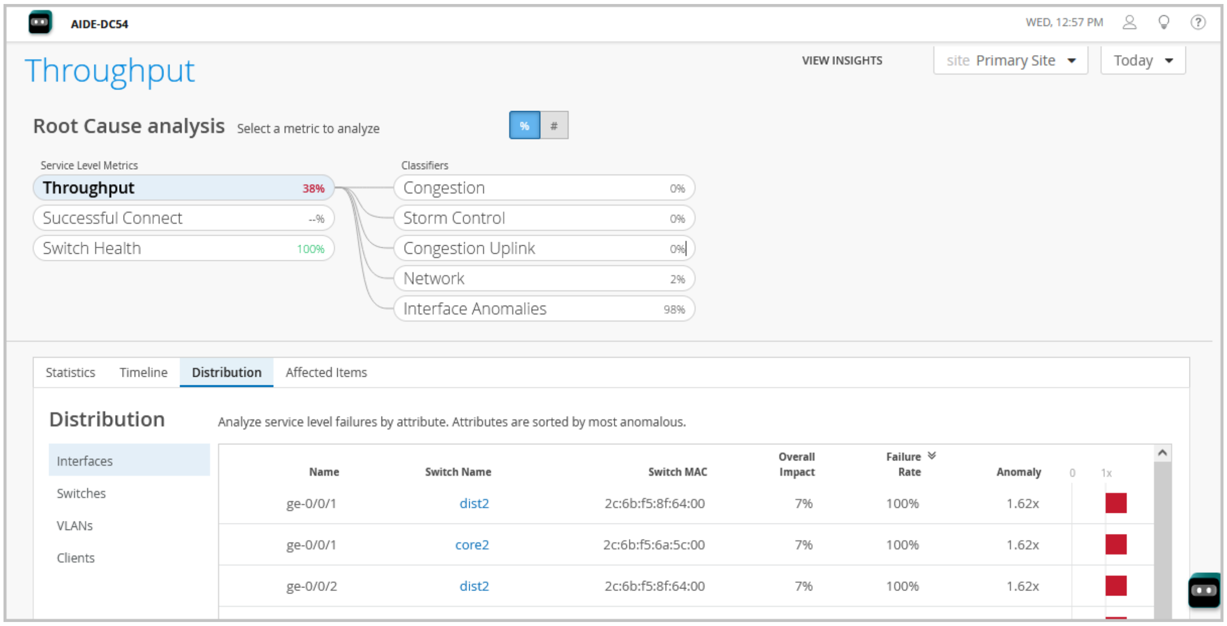

AIDE-DC54 WED, 12:57 PM

VIEW INSIGHTS site Primary Site Today

Throughput

Root Cause analysis Select a metric to analyze % #

Service Level Metrics	
Throughput	38%
Successful Connect	--%
Switch Health	100%

Classifiers	
Congestion	0%
Storm Control	0%
Congestion Uplink	0%
Network	2%
Interface Anomalies	98%

Statistics Timeline Distribution Affected Items

Distribution

Analyze service level failures by attribute. Attributes are sorted by most anomalous.

Interfaces
Switches
VLANs
Clients

Name	Switch Name	Switch MAC	Overall Impact	Failure Rate	Anomaly	0 1x
ge-0/0/1	dist2	2c:6b:f5:8f:64:00	7%	100%	1.62x	🟥
ge-0/0/1	core2	2c:6b:f5:6a:5c:00	7%	100%	1.62x	🟥
ge-0/0/2	dist2	2c:6b:f5:8f:64:00	7%	100%	1.62x	🟥

The above image indicates that the network met the throughput requirement only for 38 percentage of the time. And that the users faced throughput issues for the remaining 62 percent of the time. The classifier view shows that 98 percent of the issues that affected the throughput belonged to Interface Anomalies category, while 2 percent of issues were network issues.

To access the classifier view, click a metric (for example, Throughput) and then select a classifier (for example, Interface Anomalies). Here's a sample of the Interface Anomalies metrics view:

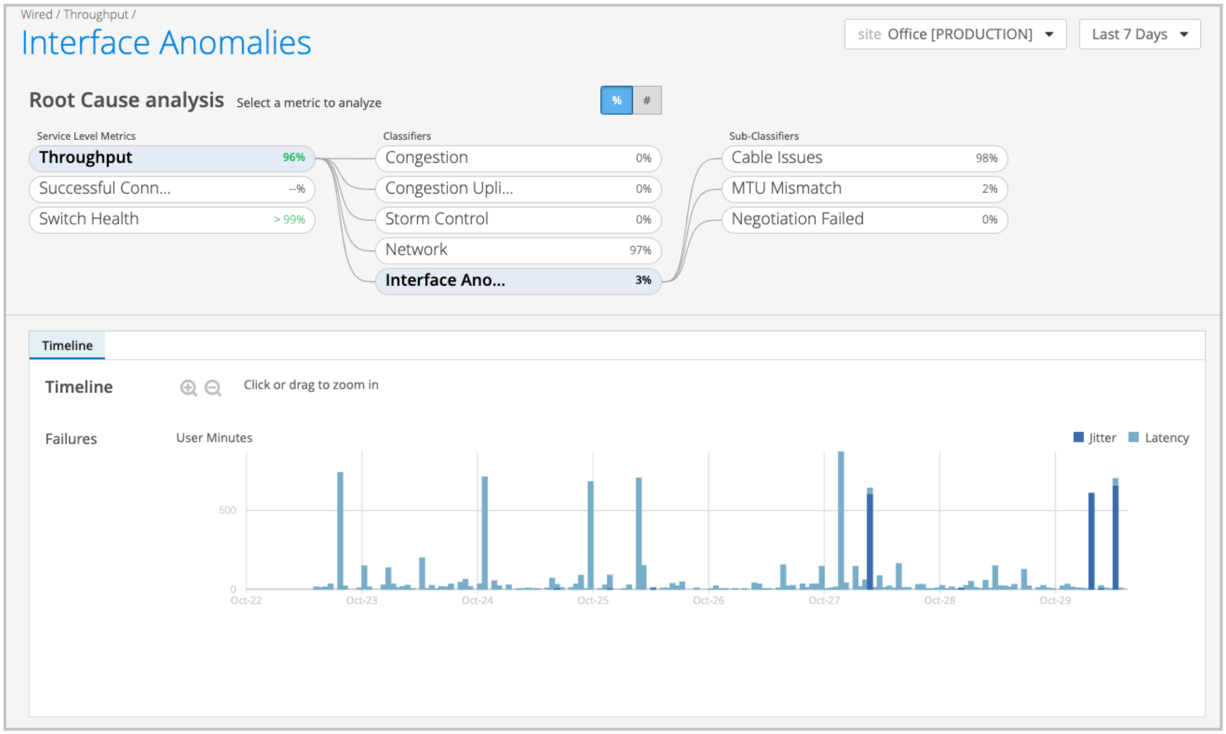

Wired / Throughput /

Interface Anomalies

site Office [PRODUCTION] Last 7 Days

Root Cause analysis Select a metric to analyze % #

Service Level Metrics	
Throughput	96%
Successful Conn...	--%
Switch Health	> 99%

Classifiers	
Congestion	0%
Congestion Upli...	0%
Storm Control	0%
Network	97%
Interface Ano...	3%

Sub-Classifiers	
Cable Issues	98%
MTU Mismatch	2%
Negotiation Failed	0%

Timeline

Timeline Click or drag to zoom in

Failures

NOTE: The classifiers do not show any data when the metric shows a success rate of 100 percent.

Throughput

The **Throughput** metric shows the percentage of the time the wired users could pass traffic without any disruptions. This classifier helps you evaluate your network and determine if it requires higher bandwidth for seamless operation. Several factors can impact network throughput, such as MTU mismatches, faulty cables, and devices negotiating at the wrong speed.

The **Throughput** SLE has five classifiers:

- **Congestion**—This classifier shows how congestion contributed to the low throughput. It counts the number of output drops resulting from congestion. When packets arrive on an interface, they are stored in a buffer. If the buffer becomes full, the device starts dropping packets (TxDrops). The classifier uses a formula that considers the following three ratios to determine if a 'bad user minute' is caused by congestion:
  - 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)
- **Congestion Uplink**—The SLE dashboard shows high congestion uplink when:
  - One of the neighbors is a switch or a router (known through LLDP).
  - The port is an STP root port.
  - The uplink port has a higher number of transmitted and received packets compared to the other ports.

  Congestion can also be caused by aggregated Ethernet links and module ports.
- **Interface Anomalies**—This classifier shows how interface anomalies contributed to the low throughput. The SLE dashboard gathers information about interface anomalies from the switches. The interface anomalies classifier is divided into the following sub-classifiers:
  - **MTU Mismatch**—As an admin, you can set a maximum transmission unit (MTU) value for each interface. The default value for Gigabit Ethernet interfaces is 1514. To support jumbo frames, you need to configure an MTU value of 9216, which is the upper limit for jumbo frames on a routed VLAN interface. It's important to ensure that the MTU value is consistent along the packet's path, as any MTU mismatch will result in discarded or fragmented packets. In Juniper switches, you can

check for MTU mismatches in the **MTU Errors** and **Input Errors** sections of the `show interface extensive` command output. Each input error or MTU error contributes to a "bad user minute" under MTU mismatch.

  - **Cable Issues**—This sub-classifier shows the user minutes affected by faulty cables in the network.
  - **Negotiation Failed**—Latency on ports can happen due to auto-negotiation failure, duplex conflicts, or user misconfiguration of device settings. Moreover, older devices may not be able to achieve maximum speed and could operate at a slower link speed of 100 Mbps. This sub-classifier identifies and helps mitigate instances of bad user time caused by these issues.

- **Storm Control**—Storm control allows the device to monitor traffic levels and drop broadcast, unknown unicast, and multicast packets when they exceed a set threshold or traffic levels. These thresholds are known as storm control levels or storm control bandwidth. By default, the storm control level is set to 80 percent of the combined broadcast, multicast, and unknown unicast traffic on all layer 2 interfaces of Juniper switches. Storm control helps prevent traffic storms, but it can also potentially throttle applications or client devices. This classifier identifies these conditions and helps users proactively mitigate throughput issues.
- **Network**—This classifier allows you to monitor user minutes when the throughput is lower than expected due to limitations in uplink capacity. It identifies issues based on the round-trip time (RTT) value of packets sent from the switch to the Mist cloud. The Network classifier has two sub-classifiers that help you locate these issues:
  - **Latency**—Displays user minutes affected by latency. The latency value is calculated based on the average value of RTT over a period of time.
  - **Jitter**—Displays user minutes affected by jitter. The jitter value is calculated by comparing the standard deviation of RTT within a small period (last 5 or 10 minutes) with the overall deviation of RTT over a longer period (day or week). You can view this information for a particular switch or site.

Switch Health

Switch health is influenced by several factors, including operating temperature, power consumption, CPU, and memory usage. Monitoring switch health is crucial because issues like high CPU usage can directly impact connected clients. For instance, if CPU utilization spikes to 100 percent, the connected APs may lose connectivity, affecting the clients' experience. The Switch Health metric identifies bad user minutes resulting from the following conditions (listed as classifiers):

- **Switch Unreachable**—The switch can't be accessed.
- **Memory**—The memory utilization is above 80 percent.
- **CPU**—The switch CPU usage is above 90 percent.

- **Temp**—The switch operating temperature exceeds the prescribed threshold range, either going above the maximum limit or below the minimum requirement. For information about the operating temperature supported by Juniper switches, refer to the switch hardware guides in Juniper documentation portal.
- **Power**—The switch power consumption is above 90 percent of the available power.

Successful Connect

The Successful Connect metric shows if a client successfully connects to the network. It helps assess the impact of connect failures and identify the issues preventing client devices from connecting to the network.

The Successful Connect metric has two classifiers:

- **Authentication**—Each time a client authenticates, a client event is generated. These could either be successful events or failure events. This classifier helps you identify issues that caused authentication failures. Here's a list of possible reasons for a dot1x authentication failure:
  - If a single switch port fails to authenticate, it could be due to a user error or misconfigured port.
  - If all switch ports fail to authenticate, it could be because:
    - The switch is not added as a NAS client in the RADIUS server.
    - There's a routing issue between the switch and the RADIUS server.
    - The RADIUS server is down.
  - If all switch ports on all switches fail to authenticate, it could indicate a temporary failure with the RADIUS server at that specific moment.
  - If a specific type of device, such as Windows devices, fails to authenticate, it may suggest an issue related to certifications.
- **DHCP**—DHCP snooping enables the switch to examine the DHCP packets and keep track of the IP-MAC address binding in the snooping table. This classifier adds a failure event every time a client connects to a network and fails to reach the ‘bound’ state within a minute.

  NOTE: The SLE dashboard shows DHCP failures only for those switches that have DHCP Snooping configured.

7

CHAPTER

Troubleshooting

Troubleshoot Your Switch

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`.

```
user@switch> ping oc-term.mistsys.net
PING ab847c3d0fcd311e9b3ae02d80612151-659eb20beaaa3ea3.elb.us-west-1.amazonaws.com
(13.56.90.212): 56 data bytes
```

   If the switch is not resolving `oc-term.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 the following table to determine which port to enable, depending on your cloud environment:

Table 8: Ports to Enable in Different Juniper Mist Clouds

Service Type	Global 01	Global 02	Europe 01
EX Switch	redirect.juniper.net (TCP 443)	redirect.juniper.net (TCP 443)	redirect.juniper.net (TCP 443)
	ztp.mist.com (TCP 443)	ztp.gc1.mist.com (TCP 443)	ztp.eu.mist.com (TCP 443)

Table 8: Ports to Enable in Different Juniper Mist Clouds *(Continued)*

Service Type	Global 01	Global 02	Europe 01
	oc-term.mistsys.net (TCP 2200)	oc-term.gc1.mist.com (TCP 2200)	oc-term.eu.mist.com (TCP 2200)

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.

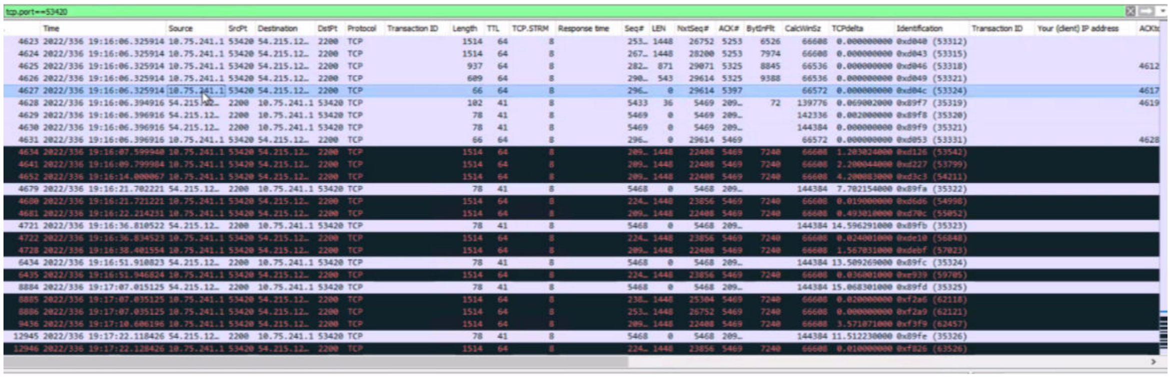

tcp.port==53420

	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)			

See also: "Packet Captures in Mist" on page 168 .

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:

Cloud-Ready LED Blink Patterns

If your switch can't connect to the cloud, LED blink patterns on the switch can help tell you why.

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 9: 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.
🟡 🟡 🟡 🟡 🟡 🟡 🟡	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.

Table 9: Cloud LED Blink Patterns *(Continued)*

CLD LEDs	Blink Pattern	Meaning
● ● ● ● ● ● ● ●	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.
● ○ ● ● ● ● ● ●	1 yellow, pause, 6 yellow	Cloud unreachable. During ZTP, the PHC could not reach the cloud.

RELATED DOCUMENTATION

No Link Title

Dynamic and Manual Packet Captures

SUMMARY

The Juniper Mist portal provides both dynamic and manual packet captures to help identify the source of communication failures between the client and AP.

IN THIS SECTION

Dynamic Packet Captures

Whenever a service-impacting event occurs between the wireless client and AP, it automatically triggers a short-term dynamic packet capture. These include DHCP issues (timeout, denied, terminated), authorization failures (RADIUS not responding, Access-Reject, incomplete authorization), and roaming issues (11r FBT and OKC authorization failures).

Captured packets are saved to the cloud, where they are associated with the triggering event in the Juniper Mist portal. You can view or download the capture from the Insights panel:

- **Monitor** > **Service Levels** | **Insights** > **Client Events**

- 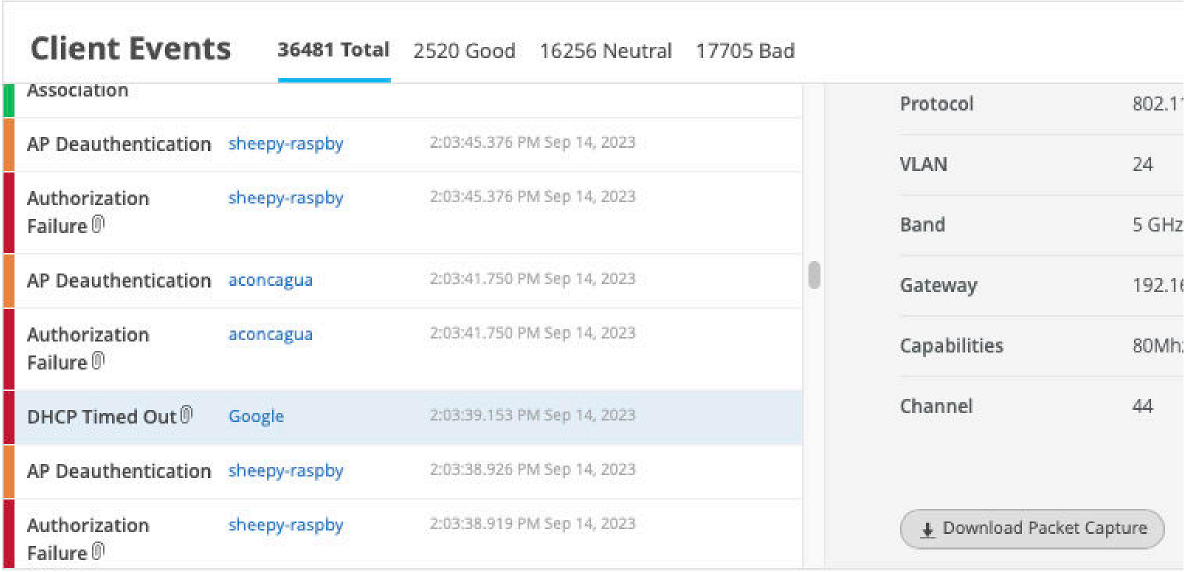

Client Events 36481 Total 2520 Good 16256 Neutral 17705 Bad

Association		
AP Deauthentication	sheepy-raspby	2:03:45.376 PM Sep 14, 2023
Authorization Failure	sheepy-raspby	2:03:45.376 PM Sep 14, 2023
AP Deauthentication	aconcagua	2:03:41.750 PM Sep 14, 2023
Authorization Failure	aconcagua	2:03:41.750 PM Sep 14, 2023
DHCP Timed Out	Google	2:03:39.153 PM Sep 14, 2023
AP Deauthentication	sheepy-raspby	2:03:38.926 PM Sep 14, 2023
Authorization Failure	sheepy-raspby	2:03:38.919 PM Sep 14, 2023

Protocol	802.1
VLAN	24
Band	5 GHz
Gateway	192.16
Capabilities	80Mh
Channel	44

Manual Packet Captures

Wired packet capture applies to the wired ports of APs (not the switch ports). WAN packet captures support SSR WAN edge device ports (not SRX WAN edge devices).

For manual packet captures, go to **Site** > **Packet Captures**, where you can

- Choose which network type to capture packets from: wired, wireless, or WAN.
- Restrict the packet capture to specific clients, WLANs, APs, or wireless bands.
- Configure the number of packets captured, packet size in bytes, and the duration of the capture session.
- Configure other capture parameters such as header inclusion and capture filters. See Table 10 on page 172 for details.

After downloading the packet capture to your computer, follow the steps below to view them in Wireshark.

Configure IEEE 802.11 on Wireshark

Packet inspection requires Wireshark. See https://www.wireshark.org for the download file and related information.

To configure Wireshark to view packets captured from the Juniper Mist portal,

1. Open the Wireshark application on your computer
2. Open the Wireshark Preferences window:

   On a Windows computer navigate to **Edit** > **Preferences**

   On a Mac computer navigate to **Wireshark** > **Preferences**
3. In the Preferences window, expand the **Protocols** menu option and scroll down to **IEEE 802.11**

   a. Select **Yes - with IV** and then click **OK**, as shown in the following image:

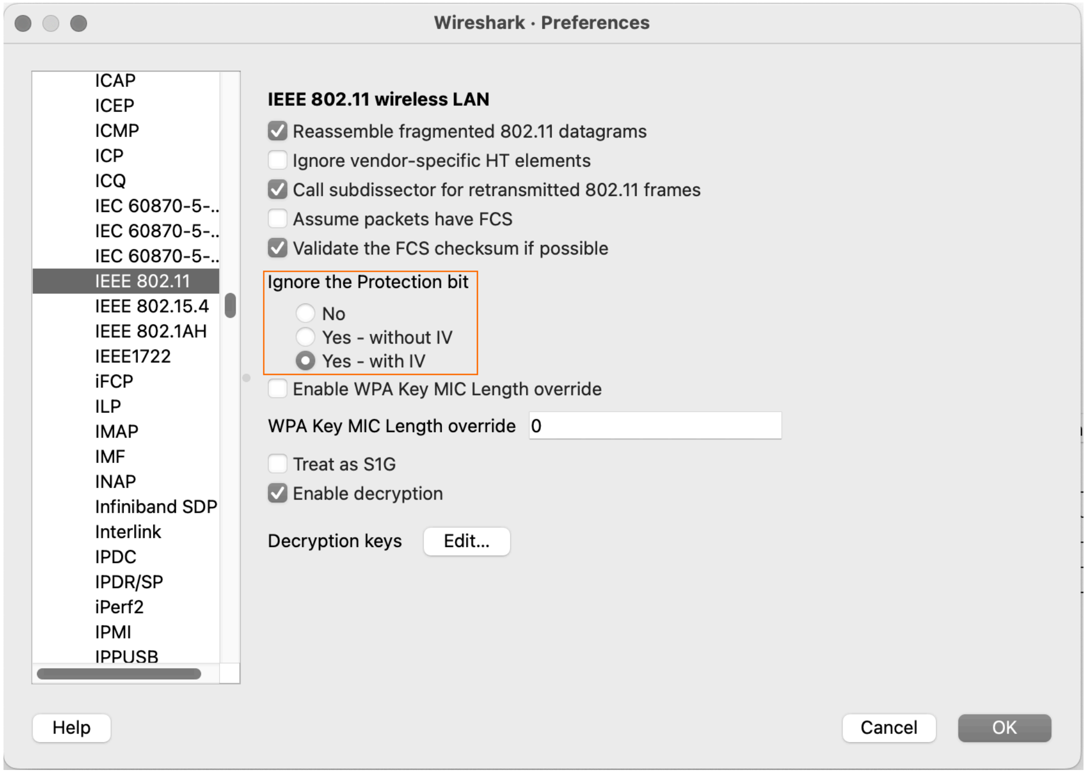

IEEE 802.11 wireless LAN

☑ Reassemble fragmented 802.11 datagrams

□ Ignore vendor-specific HT elements

☑ Call subdissector for retransmitted 802.11 frames

□ Assume packets have FCS

☑ Validate the FCS checksum if possible

Ignore the Protection bit

○ No

○ Yes - without IV

● Yes - with IV

□ Enable WPA Key MIC Length override

WPA Key MIC Length override 0

□ Treat as S1G

☑ Enable decryption

Decryption keys Edit...

View Wireless Packet Captures in Wireshark

You can capture packets from both your wired and wireless networks. The following configuration regards wireless packet, for which you can see:

- Wireless channel information
- Wireless data rate
- Received signal strength indicator (RSSI)

To accomplish this task, you must download and install the Wireshark application on your computer. In a Web browser, navigate to https://www.wireshark.org for Wireshark application downloads and detailed information about Wireshark. For additional information about Wireshark, see https://www.wireshark.org/docs/.

This topic provides minimal guidance about how to configure Wireshark for use in examining wireless packet captures gathered from the Juniper Mist portal.

1. Open the Wireshark application on your computer.
2. Open the Wireshark Preferences window:

   On a Windows computer navigate to **Edit** > **Preferences**.

   On a Mac computer navigate to **Wireshark** > **Preferences**.
3. In the Preferences window, navigate to **Appearance** > **Columns**.
4. Click the **Add (+)** button to add a new radiotap column to the Wireshark display (radiotap headers include wireless packet frames that would otherwise not be displayed. See: https://www.wireshark.org/docs/dfref/r/radiotap.html.

   Wireshark adds a new line called New Column, and the type Number.

   a. Double-click the **New Column** title and type Channel as the title.

   b. Double-click the **Type** column and select Frequency/Channel from the drop-down menu.

   c. Leave the **Displayed** column selected.
5. Repeat Step 4 two times

   a. The first time, use **Data Rate** for the column title and **IEEE 802.11 TX Rate** for the type.

   b. The second time, use **RSSI** as the column title and **IEEE 802.11 RSSI** for the type.
6. Click **OK** to save your changes.

   Wireshark will display the new columns when you open a packet capture (.pcap) file for viewing.

Manual Packet Capture Options

By default, Juniper Mist streams the packet capture session data, including beacon frames, to the Mist portal. The following table describes the packet capture options that you can use when you create a packet capture session.

Table 10: Packet Capture Options

Option Name	Option Function	Usage Notes	Firmware Notes
Include Network Headers	Include packet headers in addition to the packet data.	Packet capture works by buffering packets locally on the device, meaning there is limited space available for storage. By default, Mist truncates header data from the captured packets to reduce the size of capture files while still providing the most relevant information.	–
Local Capture	Do not stream the live capture data to the Mist GUI.	Earlier AP firmware did not support live streaming packet captures to the Juniper Mist portal.	Required for AP firmware versions before 0.10.x.
Canned Filters	Pre-defined filters that vary based on the type of packet capture you're performing.	The filters available in the list change depending on whether you're capturing wireless, wired, or WAN packets. For example, beacon frames are only available for wireless packet captures.	–
Advanced Filters	Create your own packet filters for the capture session using `tcpdump` syntax.		0.10.x or later
Expression Builder	Interactive GUI tool to build custom filters in `tcpdump` syntax for use in the capture session.	You can let the builder start the filter entry and then add to or delete from the entry manually.	0.10.x or later

Troubleshoot with Marvis

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)

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.

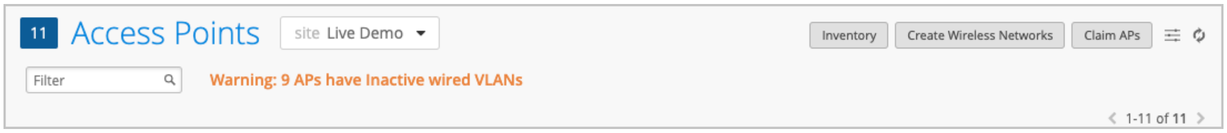

Warning: 9 APs have Inactive wired VLANs

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.

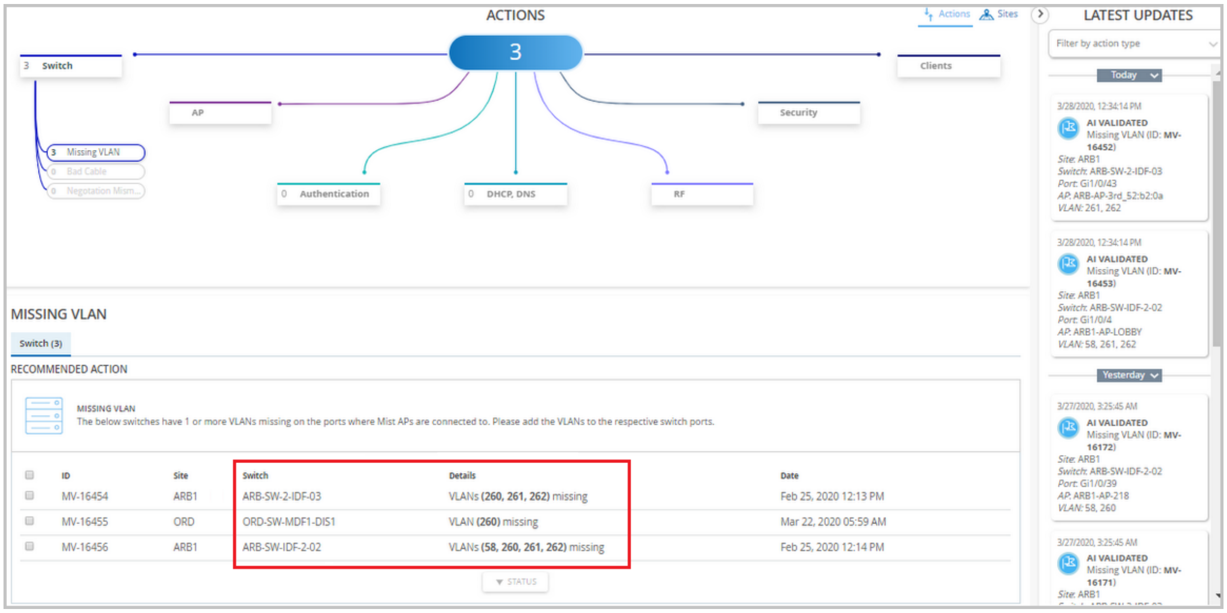

MISSING VLAN

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

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.