

# Juniper Mist Wired Assurance Configuration Guide

Published  
2025-01-15

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

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

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

*Juniper Mist Wired Assurance Configuration Guide*  
Copyright © 2025 Juniper Networks, Inc. All rights reserved.

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

## YEAR 2000 NOTICE

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

## END USER LICENSE AGREEMENT

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

# Table of Contents

About This Guide | vii

1

## Get Started

Juniper Mist Wired Assurance Overview | 2

Hardware and Software Requirements for Your Wired Network | 3

Switch Administrator Role Requirements | 3

Deploy Your Wired Network | 6

Request Help with a New Deployment of Juniper Mist Devices | 8

Explore Juniper Mist Features | 10

Port Profiles Overview | 11

Group-Based Policy Configuration Overview (Mist) | 20

Juniper CloudX Overview | 21

2

## Switch Configuration

Switch Configuration Overview (Mist) | 25

Onboard Switches to Mist Cloud | 26

Switch Onboarding Prerequisites | 26

Onboard a Greenfield Switch | 27

Onboard a Brownfield Switch | 27

Configure Switches | 30

Create a Switch Configuration Template | 31

Assign a Template to Sites | 32

Configure Switch-Specific Settings | 32

Configure Switch-Specific Settings Manually | 33

Configure Switch-Specific Settings Using the Bulk Upload Option | 33

Verify the Switch Configuration | 36

## Switch Configuration Options | 37

### Protection of Routing Engine | 76

Configure Protection of Routing Engine | 77

Verify Protection of Routing Engine Configuration | 81

### QoS Configuration | 90

### Configure SNMP on Switches | 100

Configure SNMP at the Organization Level | 101

Configure SNMP at the Site Level | 103

Configure SNMP at the Device Level | 105

### Configure DHCP Server or Relay on a Switch | 107

Prerequisites | 107

Configure DHCP Server | 107

Configure DHCP Relay | 114

### OSPF Configuration for Switches | 117

Example: Configure Basic OSPF in two EX Devices | 118

CLI Commands | 122

OSPF Events | 124

Configure OSPF at the Organization Level | 124

Configure OSPF at the Site Level | 125

### Manage or Update Configuration Settings | 125

Manage Templates Settings | 125

Update Switch Configuration Settings at the Site Level | 126

Add or Delete a CLI Configuration | 127

### Upgrade Junos OS Software on Your Switch | 128

Free Up Storage Space on Your Switch | 129

Upgrade the Junos OS Software on Your Switch | 131



[Create Recovery Snapshot for a Switch | 135](#)

[Assign a Role to Switches | 136](#)

[Locate a Switch by LED | 137](#)

[Replace a Switch | 138](#)

[Disable Remote Shell Access to Switches and Gateway Devices | 141](#)

[Connect a Switch to Mist Cloud via a Proxy Server Using Cloudx | 141](#)

[Connect a Switch to Mist Cloud via a Dynamic Proxy Server | 142](#)

[Connect a Switch to Mist Cloud via a Static Proxy Server | 142](#)

[Configure the System Log | 144](#)

### 3

## Switch Dashboards

[Switch Metrics | 148](#)

[Switch Details | 149](#)

[Switch Utilities | 161](#)

[Wired Clients | 163](#)

### 4

## Virtual Chassis Configuration

[Virtual Chassis Overview \(Juniper Mist\) | 167](#)

[Configure a Virtual Chassis Using EX2300, EX4650, or QFX5120 Switches | 171](#)

[Configure a Virtual Chassis Using EX3400, EX4100, EX4100-F, EX4300, EX4400, or EX4600 Switches | 175](#)

[Manage a Virtual Chassis Using Mist \(Add, Delete, Replace, and Modify Members\) | 180](#)

[Prerequisites | 181](#)

[Replace a Member Switch in a Virtual Chassis | 182](#)

[Replace a Non-FPC0 Member in a Virtual Chassis | 182](#)

[Replace the FPC0 Member in a Virtual Chassis | 186](#)

[Renumber the Virtual Chassis Members | 190](#)

[Reassign the Virtual Chassis Member Roles | 192](#)

Delete Virtual Chassis Members | 194

Add a Member Switch to a Virtual Chassis | 195

5

## Campus Fabric Configuration

Which Campus Fabric Topology to Choose | 200

How to Migrate a Traditional Enterprise Network to a Juniper Campus Fabric | 208

Configure Campus Fabric EVPN Multihoming | 213

Configure Campus Fabric Core-Distribution | 221

Configure Campus Fabric IP Clos | 231

6

## Wired Service Levels

Service Level Expectations (SLE) | 242

Wired SLEs Dashboard | 250

Wired Throughput SLE | 251

Wired Successful Connect SLE | 254

Switch Health SLE | 255

Switch Bandwidth SLE | 257

7

## Troubleshooting

Troubleshoot Your Switch Connectivity | 261

Troubleshooting Juniper CloudX | 267

Cloud-Ready LED Blink Patterns | 271

| Cloud-Ready Connection Process | 274

Troubleshoot with Marvis | 275

FAQs (Mist Wired) | 276

8

## Appendix

Deploy and Manage EX Series Switches at a Branch | 284

## 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 | 2](#)

[Hardware and Software Requirements for Your Wired Network | 3](#)

[Switch Administrator Role Requirements | 3](#)

[Deploy Your Wired Network | 6](#)

[Request Help with a New Deployment of Juniper Mist Devices | 8](#)

[Explore Juniper Mist Features | 10](#)

[Port Profiles Overview | 11](#)

[Group-Based Policy Configuration Overview \(Mist\) | 20](#)

[Juniper CloudX Overview | 21](#)

---

# 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 \(Mist\)](#) | 25

# Hardware and Software Requirements 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. Juniper Mist Wired Assurance supports onboarding, configuration, and management of Juniper EX Series and QFX Series switches. Click the links below to access datasheets, quick start guides, and hardware guides for these switches on the Juniper Mist Supported Hardware page:

- [EX Series Switches](#)
- [QFX Series Switches](#)



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

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

## 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	✓	×	×	×	×
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 <b>Switches</b> page or the <b>Site &gt; Switch Configuration</b> page.	✓	✓	✓	×	×
Access switch template	✓	×	×	×	×

*(Continued)*

Privileges	Super User	Network Admin (All Sites Access)	Network Admin (Site Group or Specific Sites Access)	Helpdesk	Observer
Assign switch template to sites	✓	✓	×	✓  (Applicable only to roles with access to all sites. Not available to roles with access to all site groups or specific sites.)	✓  (Applicable only to roles with access to all sites. Not available to roles with access to all site groups or specific sites.)
Enable/disable switch configuration management	✓	✓	✓	×	×
Send the switch logs to the Mist cloud	✓	✓	✓	×	×
View the Inventory page	✓	✓	×	×	✓
Assign to a site	✓	✓  (Applicable only to the site assignment option on the switch details page)	×	×	×



*(Continued)*

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

## 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	<p>Before you can configure your wired network or onboard your devices, you need to complete these tasks in the Juniper Mist™ portal:</p> <ul style="list-style-type: none"> <li>• Create your organization, set up at least one site, and activate your subscriptions.</li> <li>• 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.</li> <li>• Configure your firewall to allow Juniper Mist traffic.</li> <li>• Set up other security options as needed. For example, manage certificates, disable Juniper Mist support access, or enable Single Sign-On.</li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">Juniper Mist Quick Start</a></li> <li>• <a href="#">Firewall Configuration: Juniper Mist IP Addresses and Ports</a></li> <li>• <a href="#">Security Options</a></li> </ul>
Understand Admin Permissions	<p>Make sure that your admin account gives you the permissions that you need for your configuration tasks.</p>	<p><a href="#">"Switch Administrator Role Requirements" on page 3</a></p>
Onboard Switches	<p>Add switches to your Juniper Mist organization, either in a greenfield (new cloud-ready switches) or a brownfield (previously deployed) approach.</p>	<p><a href="#">"Onboard Switches to Mist Cloud" on page 26</a></p>

**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.	<a href="#">"Configure Switches" on page 30</a>

## Request Help with a New Deployment of Juniper Mist Devices

### SUMMARY

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

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

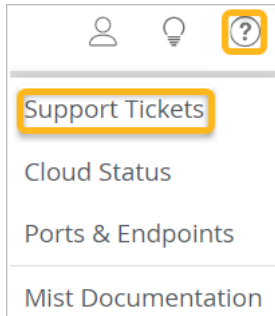


#### NOTE:

- These services do not include network design.
- Submit your request at least 48 hours in advance of your preferred appointment time.
- Available only for wireless, switching, and SD-WAN deployments.

To request help with a new deployment:

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



2. Click **Create a Ticket**.
3. For **Technology**, select the relevant technologies for your deployment:
  - **Wireless**
  - **Switching**
  - **SD-WAN**
4. For **Ticket Type**, select **Onboarding Help for New Deployment**.

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



**NOTE:** The checklist includes each technology that you selected at the top. Complete all questions in all sections that appear.

As you complete the checklist, you can use suggested hyperlinks for self-help. If you find your answers in these links, you can cancel submitting this ticket.

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



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

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

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

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

## Explore Juniper Mist Features

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 148](#)
  - ["Switch Details" on page 149](#)
  - ["Switch Utilities" on page 161](#)
- Wired Service Level Expectations (SLEs)—Use the SLE dashboards to assess the network's user experience and resolve any issues proactively. See ["Wired SLEs Dashboard" on page 250](#).
- Port Profiles—Port profiles provide a convenient way to manually or automatically provision switch interfaces. See ["Port Profiles Overview" on page 11](#).
- 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 200.

- **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 20.
- **Virtual Chassis**—The Virtual Chassis technology enables you to connect multiple individual switches together to form one logical unit and to manage the individual switches as a single unit. See ["Virtual Chassis Overview \(Juniper Mist\)"](#) on page 167.

## Port Profiles Overview

### IN THIS SECTION

- [Static Port Profiles | 12](#)
- [Dynamic Port Profiles | 12](#)
- [Best Practices in Port Configuration | 15](#)
- [System-defined Port Profiles | 17](#)

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).
- **System-defined port profiles**—By default, Juniper Mist provides you with system-defined port profiles that are preconfigured for you. These work the same way regular port profiles do, except these are available for you to use if you do not want to configure your own. The system-defined port profiles provided by Mist are as follows: ap, iot, uplink, default, and disabled.

## 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 ["Configure Switches"](#) on page 30.

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.



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.

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

See "[Configure Switches](#)" on page 30 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

```

## System-defined Port Profiles

System-defined Port Profiles are port profiles that are built into the Mist portal and are available for you to use if you do not want to configure your own port profiles. These are preconfigured for you, so there is no configuration required in order for you to be able to use them. You can, however, delete these system-defined port profiles. This functionality is only available from the Switch Templates level of configuration.

The following steps describe how to delete a system-defined port profile.



**NOTE:** The ability to delete a system-defined port profile only applies to the ap, iot, and uplink port profiles.

1. From the Mist portal, navigate to **Organization > Switch Templates**.
2. Select the appropriate Switch Template.
3. Select the system-defined port profile you wish to delete (ap, iot, or uplink).
4. Select the **trashcan** icon in the top left corner of the **Edit Port Profile** configuration.

## PORT PROFILES

Port configuration for a set of related ports

\* System defined

✖ 📄 **Edit Port Profile** ✓ ✕

Override System defined profile

Name

Port Enabled

Enabled  Disabled

Description

---

Mode

Trunk  Access

Port Network (Untagged/Native VLAN)

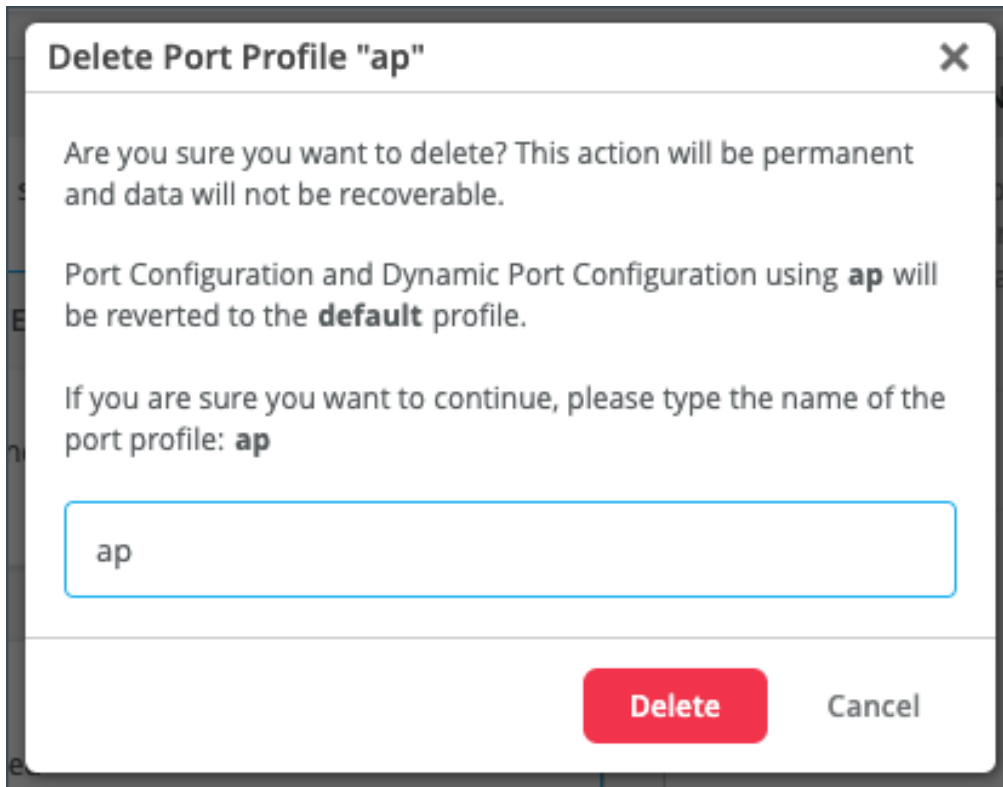
default	1
---------	---

VoIP Network

Trunk Networks

All Networks

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



**i** **NOTE:** If you delete the `ap`, `iot`, or `uplink` system-defined port profiles, any reference to these profiles at the Site or device level will revert to the default profile (port configurations or Dynamic Port profiles).

**i** **NOTE:** If you were to create your own port profile and name it "`ap`", "`iot`", or "`uplink`" (after having deleted the system-defined port profiles) it will be treated as any other user-defined port profile.

## 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:** [Campus Fabric GBP Microsegmentation](#)

---

On the Mist portal, you can configure GBP using the switch templates (**Organization > Switch Templates**), or directly from the switch configuration page (**Switches > Switch Name**). The GBP configuration involves creating GBP tags and including them in switch policies. The GBP tags enable you to group users and resources. In a GBP, you match a user group tag to a resource group tag to provide the specified users access to the specified resources.

The following video takes you through the steps involved in configuring a GBP:



**Video:** [Group Based Policy Demo](#)


---

See also: [Microsegmentation with GBP Using Mist Wired Assurance](#).

## Juniper CloudX Overview

Juniper CloudX, integrated natively into Junos OS, is an advanced architecture that ensures faster and secure communication between Juniper switches and the Mist cloud. It is responsible for creating a secure connection between the switch and the Mist cloud. CloudX-enabled switches can be monitored and managed by cloud services.





CloudX applies to both new and existing switches. It enables the new switches to communicate directly over HTTPS 443 when they are onboarded to Mist cloud via ZTP. With CloudX enabled, the existing switches that are connected to the Mist cloud via TCP port 2200 will have their connection switched to CloudX with no impact on the data plane. For switches to connect and communicate using CloudX over TCP 443, the following [firewall ports](#) must be opened: `jma-terminator.[xx].mist.com`(TCP 443). The variable [xx] should be replaced by the environment name.

#### Benefits of CloudX:

- Keeps the data on the cloud up to date. Events are sent to the cloud every 10-15 seconds and stats are updated every 60 seconds.
- Leverages the Junos Telemetry Interface (JTI), which ensures asynchronous and faster communication by bypassing any polling from the cloud to the switch.
- Enables switches to connect to the cloud over HTTPS port 443, like Mist APs. You do not need to open any non-standard ports on the firewall.
- Enables switches to communicate with the Mist cloud via a proxy server. You can statically define a proxy server or dynamically send proxy server details via DHCP Option 43. For more information, see "[Connect a Switch to Mist Cloud via a Proxy Server Using Cloudx](#)" on page 141.
- Offers packet capture for switches on the Mist Cloud. You can initiate packet capture on a single switch port or a range of ports. You can leverage the on-demand packet capture feature in Mist to view transit traffic or control traffic.

#### Availability of CloudX

The following table lists the platforms that support CloudX in different Junos releases. The table lists multiple Junos versions for each platform. We recommend that you upgrade the switch to a [Junos suggested release](#) for the CloudX support.

**Table 2: CloudX-Supported Platforms**

Platforms	Supported Junos Release	Availability
EX2300/EX3400	21.4R3-S4 and above 22.4R2-S1 and above 22.4R3 and above 23.4R2 and above	Beta
EX4400/EX4100	22.4R2-S1 and above 22.4R3 and above 23.4R2 and above	Generally Available (except in the Global 01 cloud instance)
EX4650/QFX5120	23.4R2 and above	Beta
Note that the following Junos versions do not support CloudX: 23.1R1 and 22.1-22.3.		

The EX4100 and EX4400 switches running the supported Junos versions have CloudX enabled across all Mist cloud environments. Support for the other platforms listed in the above table is in beta state. If you want to enable the beta version of CloudX on a supported platform (EX2300, EX3400, EX4650, and QFX5120), contact Juniper support.

If you face any issues with CloudX, you can troubleshoot it by following the steps listed in ["Troubleshooting Juniper CloudX" on page 267](#).

# 2

CHAPTER

## Switch Configuration

---

- Switch Configuration Overview (Mist) | 25
  - Onboard Switches to Mist Cloud | 26
  - Configure Switches | 30
  - Switch Configuration Options | 37
  - Protection of Routing Engine | 76
  - QoS Configuration | 90
  - Configure SNMP on Switches | 100
  - Configure DHCP Server or Relay on a Switch | 107
  - OSPF Configuration for Switches | 117
  - Manage or Update Configuration Settings | 125
  - Upgrade Junos OS Software on Your Switch | 128
  - Create Recovery Snapshot for a Switch | 135
  - Assign a Role to Switches | 136
  - Locate a Switch by LED | 137
  - Replace a Switch | 138
  - Disable Remote Shell Access to Switches and Gateway Devices | 141
  - Connect a Switch to Mist Cloud via a Proxy Server Using Cloudx | 141
  - Configure the System Log | 144
-

# 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 31](#) 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 11](#).

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

- [Switch Onboarding Prerequisites | 26](#)
- [Onboard a Greenfield Switch | 27](#)
- [Onboard a Brownfield Switch | 27](#)



**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 you want to onboard multiple cloud-ready switches together, you can do that via the Juniper Mist portal, by using the activation code associated with the purchase order.

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

For a quick demo, watch the following video:



Video: [Onboard One or More Switches Using a Web Browser](#)

## Onboard a Brownfield Switch

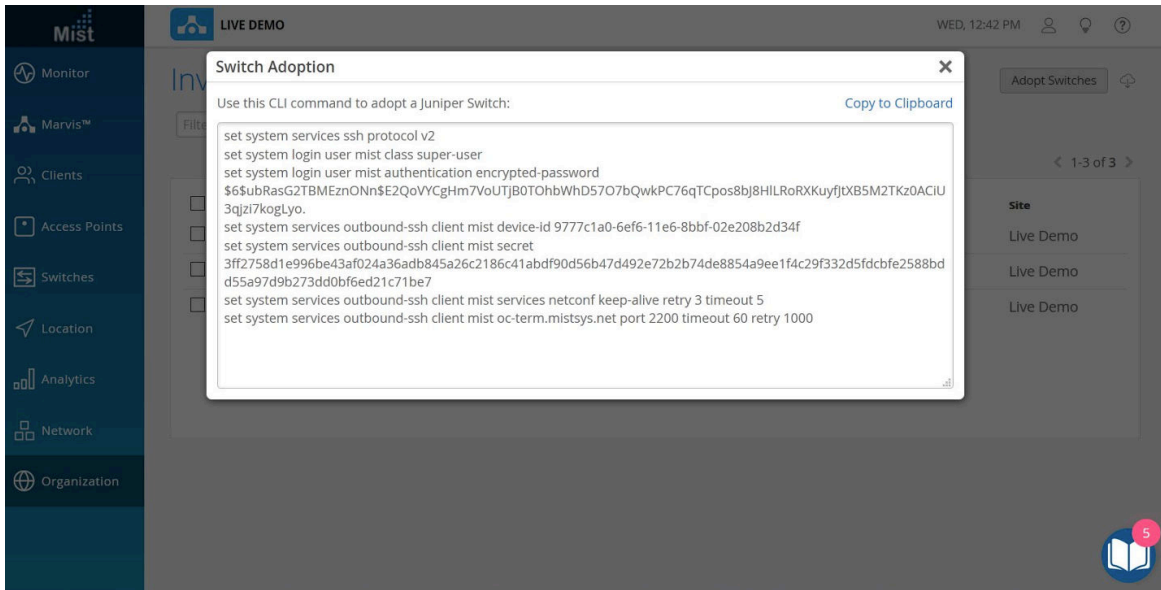
It is important to back up your existing Junos OS configuration on the switch before activating a brownfield switch because when the switch is adopted for management from the Juniper Mist cloud, the old configuration is replaced. Back up your existing Junos OS configuration by running the `request system configuration rescue save` 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 an SSH connection to the Juniper Mist cloud over TCP port 2200 (the switch connection is initiated from a management interface and is used for setting up configuration and sending telemetry data).

Figure 1: The Switch Adoption Page



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**.



For a quick demo, watch the following video:



Video: [Onboard a Brownfield Switch](#)

## Configure Switches

### IN THIS SECTION

- [Create a Switch Configuration Template | 31](#)
- [Assign a Template to Sites | 32](#)
- [Configure Switch-Specific Settings | 32](#)
- [Verify the Switch Configuration | 36](#)

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



**NOTE:** If you prefer, you can import the template settings in a JSON file instead of manually entering the information. To import the settings, click **Import Template**. To get a JSON file with the configuration settings that can be customized and imported, open an existing configuration template of your choice and click **Export**. For more information, refer to ["Manage Templates Settings" on page 125](#).

3. Enter your settings in these sections:
  - ["All Switches" on page 38](#)
  - ["Management" on page 53](#)
  - ["Shared Elements" on page 56](#)
  - ["Select Switches Configuration" on page 67](#)
  - ["Switch Policy Labels \(Beta\)" on page 73](#)
  - ["Switch Policy \(Beta\)" on page 75](#)

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

The **Confirm changes** window appears.

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

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

## Assign a Template to Sites

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

To assign a template to one or multiple sites:

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

The **Switch Templates** page appears.

2. Click the template that you want to assign to sites.

The **Switch Templates: *Template-Name*** page appears.

3. Click **Assign to Sites**.

The **Assign Template to Sites** window appears.

4. Select the sites to which you want to apply the template and then click **Apply**.

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

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

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

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

## Configure Switch-Specific Settings


### IN THIS SECTION

- [Configure Switch-Specific Settings Manually | 33](#)
- [Configure Switch-Specific Settings Using the Bulk Upload Option | 33](#)

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:
    - **INFO**— Configure the details on the INFO tab. The details include a hostname for the switch and the role of the switch in the network (example: Access).
    - **IP Configuration (Out of Band)**—The management interface settings. If you want to override these settings, select the **Override Site/Template Settings** check box, and then make the changes. You can specify if the IP address is static or DHCP-based. You can also enable or disable **Dedicated Management VRF** (out of band). For all standalone devices or Virtual Chassis running Junos version 21.4 or later, this feature confines the management interface to non-default virtual routing and forwarding (VRF) instances. Management traffic no longer has to share a routing table with other control traffic or protocol traffic.
    - **IP Configuration**—The IRB interface settings for inter-VLAN routing. If you want to override these settings, select the **Override Site/Template Settings** check box, and then make the changes. You can specify if the IP address is static or DHCP-based. You can specify multiple IP addresses in the Additional IP Configuration section. Mist supports IPv4 and IPv6 addresses. The IPv6 support is available for IP address and subnet mask.
-  **NOTE:** If the IP address specified in the Additional IP Configuration section under IP Configuration does not fall within the scope of the subnet configured in the associated network (VLAN), the IP Configuration window displays a warning message to indicate the mismatch.
- **Port Configuration**—Configure switch ports and apply port profiles to them. L3 interface in Port Configuration support IPv4 and IPv6 addresses. The IPv6 support is available for IP address and subnet mask.
4. If you want to override the template settings applied to the switch, follow these steps:
    - a. Select the **Override Site/Template Settings** check box in relevant tiles.
    - b. Edit the settings and then click **Save**. The changes are immediately applied to the switch.

## Configure Switch-Specific Settings Using the Bulk Upload Option

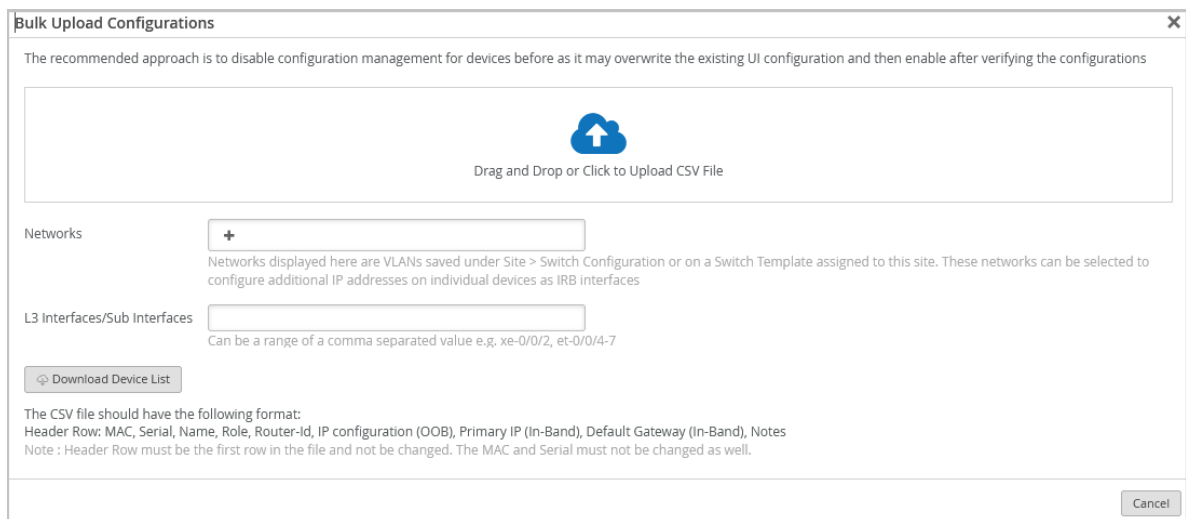
If you don't want to manually configure the switch-specific settings on each switch, you can configure the settings by uploading them through a CSV file. You can upload the settings for one or more switches at once. You can upload the following settings: MAC address, serial number, switch name, switch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

To upload the switch level settings:

1. Click **Switches**.
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.



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,

MAC	Serial	Name	Role	Router-Id	IP configuration (OOB)	Primary IP (In-Band)	Default Gateway (In-Band)	No
xxxxxxxxxxxx	serial001	sample_hostname	sample_role	x.x.x.x	x.x.x.x/y or {{siteVar}},x.x	x.x.x.x/y or {{siteVar}},x.x	x.x.x.x	
8c:01:1a:b0:10:0c	FJ0223AV0972	bhanu-VC	test-3	33.3.3.3	dhcp	dhcp:default		s
8c:01:1a:b0:60:0c	FJ0223AV0353	wing_A			dhcp	dhcp:default		
8c:01:1a:b0:00:0c	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 149 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.

The screenshot displays the Mist portal interface for a switch named 'ld-cup-idf-d-sw1\_1'. The 'Front Panel' tab is active, showing a grid of ports. A tooltip is open for port 'ge-0/0/45', which is connected to a 'Mist AP'. The tooltip provides the following configuration and status information:

- 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...01
- IP Address: 19...
- 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 156 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 125.

## SEE ALSO

[Manage or Update Configuration Settings | 125](#)

[Switch Details | 149](#)

[Wired SLEs Dashboard | 250](#)

# Switch Configuration Options

## SUMMARY

Use this information to configure your switches.

## IN THIS SECTION

- [Overview | 38](#)
- [All Switches | 38](#)
- [Management | 53](#)
- [Shared Elements | 56](#)
- [Shared Elements—Port Profiles | 60](#)
- [Select Switches Configuration | 67](#)
- [Switch Policy Labels \(Beta\) | 73](#)
- [Switch Policy \(Beta\) | 75](#)



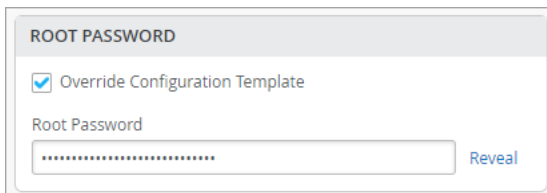
## Overview

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

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

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

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




**NOTE:** The fields that support configuration through site variable have a help text showing the site variable configuration format underneath them. To configure site variables, follow the steps provided in [Configure Site Variables](#). For more information about the switch configuration process and switch templates, see "[Configure Switches](#)" on page 30.

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

## All Switches

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

### All Switches Configuration

#### AUTHENTICATION SERVERS

Authentication Servers

Radius

Authentication Servers

No servers defined

[Add Server](#)

Timeout:  (0 - 1000 seconds)

Retries:  (0 - 100)

Enhanced Timers **i**

Enabled  Disabled

Load Balance **i**

Enabled  Disabled

Accounting Servers

No servers defined

[Add Server](#)

Interim Interval:  (0 - 3600 seconds)

#### TACACS+

Enabled  Disabled

#### CLI CONFIGURATION

Additional CLI Commands **i**

```
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

[Add Port Mirror](#)

Table 3: All Switches Configuration Options

Field	Description
RADIUS	<p>Choose an authentication server for validating usernames and passwords, certificates, or other authentication factors provided by users.</p> <ul style="list-style-type: none"> <li> <b>Mist Auth</b>—Configure Juniper Mist Access Assurance, a cloud-based authentication service, on your switch. For this option to work, you must use a port with dot1x or MAB authentication. For information, see the <a href="#">Juniper Mist Access Assurance Guide</a>. </li> <li> <b>RADIUS</b>—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. </li> </ul> <p>The default port numbers are:</p> <ul style="list-style-type: none"> <li>port 1812 for the authentication server</li> <li>port 1813 for the accounting server</li> </ul> <p><b>NOTE:</b> If you want to set up RADIUS authentication for Switch Management access (for the switch CLI login), you need to include the following CLI commands in the Additional CLI Commands section in the template:</p> <pre> set system authentication-order radius set system radius-server <i>radius-server-IP</i> port 1812 set system radius-server <i>radius-server-IP</i> secret <i>secret-code</i> set system radius-server <i>radius-server-IP</i> source-address <i>radius-Source-IP</i> set system login user remote class <i>class</i> </pre> <p>For RADIUS or TACACS+ local authentication to the Switch, it is necessary to create a remote user</p>

Table 3: All Switches Configuration Options (Continued)

Field	Description
	<p>account or a different login class. To use different login classes for different RADIUS-authenticated users, create multiple user templates in the Junos OS configuration by using the following CLI commands in the Additional CLI Commands section:</p> <pre> set system login user R0 class read-only set system login user OP class operator set system login user SU class super-user set system login user remote full-name "default remote access user template" set system login user remote class read-only </pre>
TACACS+	<p>Enable TACACS+ for centralized user authentication on network devices.</p> <p>To use TACACS+ authentication on the device, you must configure information about one or more TACACS+ servers on the network. You can also configure TACACS+ accounting on the device to collect statistical data about the users logging in to or out of a LAN and send the data to a TACACS+ accounting server.</p> <p>In addition, you can specify a user role for TACACS+ authenticated users within switch configuration. The following user roles are available: None, Admin, Read, Helpdesk. When the TACACS+ authenticated users do not have a user account configured on the local device, Junos assigns them a user account named 'remote' by default.</p> <p>The port range supported for TACACS+ and accounting servers is 1 to 65535.</p> <p><b>NOTE:</b> For TACACS+ to authenticate into the Switch, a similar login user as defined in the RADIUS section above needs to be created.</p>

**Table 3: All Switches Configuration Options (Continued)**

Field	Description
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 3: All Switches Configuration Options (Continued)

Field	Description
SNMP	<p>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:</p> <ul style="list-style-type: none"> <li>• <b>Options under SNMPv2 (V2)</b> <ul style="list-style-type: none"> <li>• <b>General</b>—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.</li> <li>• <b>Client</b>—Define a list of SNMP clients. You can add multiple client lists. This configuration includes a name for the client list and IP addresses of the clients (in comma separated format). Each client list can have multiple clients. A client is a prefix with /32 mask.</li> <li>• <b>Trap Group</b>—Create a named group of hosts to receive the specified trap notifications. At least one trap group must be configured for SNMP traps to be sent. The configuration includes the following fields: <ul style="list-style-type: none"> <li>• <b>Group Name</b>—Specify a name for the trap group.</li> <li>• <b>Categories</b>—Choose from the following list of categories. You can select multiple values. <ul style="list-style-type: none"> <li>• authentication</li> </ul> </li> </ul> </li> </ul> </li> </ul>

Table 3: All Switches Configuration Options (*Continued*)

Field	Description
	<ul style="list-style-type: none"> <li>• chassis</li> <li>• configuration</li> <li>• link</li> <li>• remote-operations</li> <li>• routing</li> <li>• services</li> <li>• startup</li> <li>• vrrp-events</li> <li>• Targets—Specify the target IP addresses. You can specify multiple targets.</li> <li>• Version—Specify the version number of SNMP traps.</li> <li>• <b>Community</b>—Define an SNMP community. An SNMP community is used to authorize SNMP clients by their source IP address. It also determines the accessibility and permissions (read-only or read-write) for specific MIB objects defined in a view. You can include a client list, authorization information, and a view in the community configuration.</li> <li>• <b>View</b>(Applicable to both SNMPv2 and SNMPv3)—Define a MIB view to identify a group of MIB objects. Each object in the view shares a common object identifier (OID) prefix. MIB views allow an agent to have more control over access to specific branches and objects within its MIB tree. A view is made up of a name and a collection of SNMP OIDs, which can be explicitly included or excluded.</li> <li>• <b>Options under SNMPv3 (V3)</b></li> </ul>

Table 3: All Switches Configuration Options (Continued)

Field	Description
	<ul style="list-style-type: none"> <li data-bbox="878 344 1403 548">• <b>General</b>—Specify the system's name, location, administrative contact information, and a brief description of the managed system. When using SNMPv2, configure an engine ID, which serves as a unique identifier for SNMPv3 entities.</li> <li data-bbox="878 590 1414 1079">• <b>USM</b>—Configure the user-based security model (USM) settings. This configuration includes a username, authentication type, and an encryption type. You can configure a local engine or a remote engine for USM. If you select a remote engine, specify an engine identifier in hexadecimal format. This ID is used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host. If you specify the Local Engine option, the engine ID specified on the General tab is considered. If no engine ID is specified, <b>local mist</b> is configured as the default value.</li> <li data-bbox="878 1121 1414 1577">• <b>VACM</b>—Define a view-based access control model (VACM). A VACM lets you set access privileges for a group. You can control access by filtering the MIB objects available for read, write, and notify operations using a predefined view (you must define the required views first from the Views tab). Each view can be associated with a specific security model (v1, v2c, or usm) and security level (authenticated, privacy, or none). You can also apply security settings (you have the option to use already defined USM settings here) to the access group from the Security to Group settings.</li> <li data-bbox="878 1619 1386 1709">• <b>Notify</b>— Select SNMPv3 management targets for notifications, and specify the notification type. To configure this, assign a name to the</li> </ul>



Table 3: All Switches Configuration Options (Continued)

Field	Description
	<p>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.</p> <ul style="list-style-type: none"> <li>• <b>Target</b>—Configure the message processing and security parameters for sending notifications to a particular management target. You can also specify the target IP address here.</li> <li>• <b>View</b>(Applicable to both SNMPv2 and SNMPv3)—Define a MIB view to identify a group of MIB objects. Each object in the view shares a common object identifier (OID) prefix. MIB views allow an agent to have more control over access to specific branches and objects within its MIB tree. A view is made up of a name and a collection of SNMP OIDs, which can be explicitly included or excluded.</li> </ul> <p>For more information, see "<a href="#">Configure SNMP on Switches</a>" on page 100.</p>

Table 3: All Switches Configuration Options (Continued)

Field	Description
STATIC ROUTE	<p>Configure static routes. The switch uses static routes when:</p> <ul style="list-style-type: none"> <li>• It doesn't have a route with a better (lower) preference value.</li> <li>• It can't determine the route to a destination.</li> <li>• It needs to forward packets that can't be routed.</li> </ul> <p>Mist supports IPv4 and IPv6 addresses for static routes. The IPv6 support is available for destination and next hop addresses.</p> <p>Types of static routes supported:</p> <ul style="list-style-type: none"> <li>• <b>Subnet</b>—If you select this option, specify the IP addresses for the destination network and the next hop.</li> <li>• <b>Network</b>—If you select this option, specify a VLAN (containing a VLAN ID and a subnet) and the next hop IP address.</li> <li>• <b>Metric</b>—The metric value for the static route. This value helps determine the best route among multiple routes to a destination. Range: 0 to 4294967295.</li> <li>• <b>Preference</b>—The preference value is used to select routes to destinations in external autonomous systems (ASs) or routing domains. Routes within an AS are selected by the IGP and are based on that protocol's metric or cost value. Range: 0 to 4294967295.</li> <li>• <b>Discard</b>—If you select this check box, packets addressed to this destination are dropped. Discard takes precedence over other parameters.</li> </ul> <p>After specifying the details, click the check mark (✓) on the upper right of the <b>Add Static Route</b> window to add the configuration to the template.</p>

Table 3: All Switches Configuration Options (Continued)

Field	Description
CLI CONFIGURATION	<p>To configure any additional settings that are not available in the template's GUI, you can use <b>set</b> CLI commands.</p> <p>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:</p> <pre>set system login message \n\n Warning! This switch is managed by Mist. Do not make any CLI changes.</pre> <p>To delete a CLI command that was already added, use the delete command, as shown in the following example:</p> <pre>delete system login message \n\n Warning! This switch is managed by Mist. Do not make any CLI changes.</pre> <p><b>NOTE:</b> Ensure that you enter the complete CLI command for the configuration to be successful.</p>
OSPF	<p>From this tile, you can:</p> <ul style="list-style-type: none"> <li>• Define an Open Shortest Path First (OSPF) area. OSPF is a link-state routing protocol used to determine the best path for forwarding IP packets within an IP network. OSPF divides a network into areas to improve scalability and control the flow of routing information. For more information about OSPF areas, see this Junos documentation: <a href="#">Configuring OSPF Areas</a>.</li> <li>• Enable or disable OSPF configuration on the switch.</li> </ul>

Table 3: All Switches Configuration Options (Continued)

Field	Description
DHCP SNOOPING	<p>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 DHCP OFFER packets on untrusted ports, assuming they originate from unauthorized DHCP servers.</p> <p>DHCP configuration has the following options:</p> <ul style="list-style-type: none"> <li>• All Networks— Select the All Networks check box to enable DHCP snooping on all VLANs.</li> <li>• Networks—If you want to enable DHCP snooping only on specific networks, click Add (+) in the Networks box and add the required VLANs.</li> <li>• 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.</li> </ul> <p>You can check ARP statistics by using the following CLI commands: <code>show dhcp-security arp inspection statistics</code>, and <code>show log messages   match DAI</code>.</p> <p>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.</p> <ul style="list-style-type: none"> <li>• 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 not have matching entries in the database, IP Source Guard discards the packet.</li> </ul>

Table 3: All Switches Configuration Options (Continued)

Field	Description
	<p><b>NOTE:</b> IP Source Guard works only with single-suplicant 802.1X user authentication mode.</p> <p><b>NOTE:</b></p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• You need to enable VLAN on the switch for the DHCP snooping configuration to take effect. So, you need to create a port profile (described later in this document) with the VLAN included and apply the profile to the switch port.</li> </ul> <p>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 <code>show dhcp-security binding</code>. This command will provide you with information about the DHCP bindings recorded in the snooping database.</p> <p>For more information, see <a href="#">DHCP Snooping and Port Security Considerations</a>.</p> <p><b>NOTE:</b> You need to enable this feature if you want to view the <b>DHCP</b> issues for the switch under the <b>Successful Connect</b> SLE metric.</p>
SYSLOG	<p>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.</p> <p>For help with the configuration options, see "<a href="#">Configure the System Log</a>" on page 144.</p>

Table 3: All Switches Configuration Options (Continued)

Field	Description
PORT MIRRORING	<p data-bbox="841 352 1101 380">Configure port mirroring.</p> <p data-bbox="841 415 1398 548">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:</p> <ul data-bbox="841 583 1398 1031" style="list-style-type: none"><li data-bbox="841 583 1398 863">• <b>Input:</b> 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.</li><li data-bbox="841 898 1398 1031">• <b>Output:</b> 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.</li></ul>

Table 3: All Switches Configuration Options (Continued)

Field	Description
Routing Policy	<p>Configure routing policies for the entire organization (Organization &gt; Switch Templates) or for a site (Site &gt; Switch Configuration). These routing policies will only be pushed to the switch configuration if it is tied to the BGP Routing Protocol. The Routing policies that are already defined inside the BGP tab of a switch will now appear on the Routing Policy tab. The routing policies are tied to protocols such as BGP or OSPF. A routing policy framework is composed of default rules for each routing protocol. These rules determine which routes the protocol places in the routing table and advertises from the routing table. Configuration of a routing policy involves defining terms, which consist of match conditions and actions to apply to matching routes.</p> <p>To configure a routing policy:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add Routing Policy</b> on the Routing Policy tile.</li> <li>2. Provide a name to the policy, and then click <b>Add Terms</b>.</li> <li>3. Provide a name to the term and specify other match details such as: <ul style="list-style-type: none"> <li>• Prefix</li> <li>• AS Path</li> <li>• Protocol</li> <li>• Community—A route attribute used by BGP to administratively group routes with similar properties.</li> <li>• Then—Then action (Accept or Reject) to be applied on the matching routes.</li> <li>• Add Action—Additional actions such as prepend AS path, set community, and set local preference.</li> </ul> </li> </ol>

Table 3: All Switches Configuration Options (*Continued*)

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

## Management

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

The screenshot displays the Management section of the Switch Configuration page, organized into several panels:

- CONFIGURATION REVERT TIMER**: A panel with a title and an information icon. It contains a text input field labeled "Configuration Revert Timer" with the value "10".
- ROOT PASSWORD**: A panel with a title and an information icon. It contains a text input field labeled "Root Password" and a "Reveal" button.
- PROTECTION OF ROUTING ENGINE**: A panel with a title and an information icon. It contains a label "Protect Routing Engine" and two radio buttons: "Enabled" (unselected) and "Disabled" (selected).
- LOCAL USERS**: A panel with a title and an "Add User" button. It contains a table with two columns: "Username" and "Class". The table is currently empty.
- LOGIN BANNER**: A panel with a title and an information icon. It contains a text area labeled "Login Banner" with a small icon to its left.
- IDLE TIMEOUT**: A panel with a title and an information icon. It contains a text input field labeled "Idle Timeout".



**Table 4: Management Configuration Options**

Option	Notes
Configuration Revert Timer	This feature helps restore connectivity between a switch and the Mist cloud if a configuration change causes the switch to lose connection. It automatically reverts the changes made by a user and reconnects to the cloud within a specified time duration. By default, this time duration is set to 10 minutes for EX Series switches. You can specify a different time duration.
Root Password	A plain-text password for the root-level user (whose username is root).

Table 4: Management Configuration Options *(Continued)*

Option	Notes
Protection of Routing Engine	<p>Enable this feature to ensure that the Routing Engine accepts traffic only from trusted systems. This configuration creates a stateless firewall filter that discards all traffic destined for the Routing Engine, except packets from specified trusted sources. Protecting the Routing Engine involves filtering incoming traffic on the router's lo0 interface. Enabling Protection of Routing Engine on Juniper Switches is suggested as the best practice.</p> <p>When Protection of Routing Engine is enabled, Mist by default ensures that the following services (if configured) are allowed to communicate with the switch: BGP, BFD, NTP, DNS, SNMP, TACACS, and RADIUS.</p> <p>If you need additional services that need access to the switch, you can use the Trusted Networks or Services section. If you want to set up access to the switch via ssh, select the ssh option under Trusted Services. If you need to allow switch to respond to pings, select the icmp option under Trusted Services.</p> <p>If you have other segments that you would like to reach the switch from, you can add them under Trusted Networks or Trusted IP/Port/Protocol.</p> <p>For more information, refer to <a href="#">Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources</a> and <a href="#">Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods</a>.</p>
Local Users	<p>Create a local user account on the switch for device management purposes. To create a user account, click <b>Add User</b> and then define a username, login class (Operator, Read-only, Super User, or Unauthorized), and a password.</p>
Idle Timeout	<p>The maximum number of minutes that a remote shell session can be idle. When this limit is reached, users are logged out. (Valid Range: 1-60).</p>

Table 4: Management Configuration Options (*Continued*)

Option	Notes
Login Banner	Enter text that you want users to see when they log in to the switch. Example: "Warning! This switch is managed by Juniper Mist. Do not make any CLI changes." You can enter up to 2048 characters.

## Shared Elements

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

The screenshot displays the 'Shared Elements' configuration interface, which is organized into three main panels:

- NETWORKS:** This panel lists 'Named VLAN IDs that can be used by Port Profiles'. It includes a search bar and an 'Add Network' button. The list contains:
  - \* camera\_network (30)
  - \* corp\_network (40)
  - \* default (1)
  - \* iot\_network (20)
- PORT PROFILES:** This panel allows for 'Allow System Defined Profiles (ap, iot, and uplink)' to be enabled or disabled. It lists 'Profiles' that are 'Template Defined':
  - \* ap (default(1), trunk, edge)
  - \* camera\_device (camera\_n... (30), access, edge)
  - \* corp\_device (corp\_netw... (40), access, edge)
  - \* default (default(1), access)
  - \* disabled (port disabled)
  - \* iot (default(1), access, edge)
  - \* iot\_device (iot\_network(20), access, edge)
  - \* mist\_ap (default(1), trunk, edge)
  - \* restricted\_device (restricted(99), access, edge)
- DYNAMIC PORT CONFIGURATION:** This panel is used to 'Apply port profiles to ports based on properties of connected clients'. It includes an 'Override Configuration Template' checkbox and a list of rules:
  - \* MAC: ec:3ef7:c6:80:84 (corp\_device)
  - \* LLDP Chassis ID: d4:20:b0 (mist\_ap)

Table 5: Shared Elements Configuration Options

Option	Notes
Networks	<p>Add or update VLANs, which you can then use in your port profiles.</p> <p>For each VLAN, enter the name, VLAN ID, and subnet. You can specify IPv4 or IPv6 address for the subnet. See the on-screen information for more tips.</p> <p>On this tile, you have an option to hide the networks that are not used in a user-defined port profiles or L3 sub-interfaces. This feature helps you quickly identify those networks that are in use and those that are not in use.</p>
Port Profiles	<p>Add or update port profiles. For help with the profile options, see the on-screen tips and <a href="#">"Shared Elements–Port Profiles" on page 60</a>.</p> <p>On this tile, you have an option to hide the port profiles that are not used in any static or dynamic port configurations defined by users. This feature helps you quickly identify those port profiles that are in use and those that are not in use.</p>

Table 5: Shared Elements Configuration Options *(Continued)*

Option	Notes
Dynamic Port Configuration	<p data-bbox="837 352 1403 485">Dynamic port profiling uses a set of device properties of the connected client device to automatically associate pre-configured port and network settings to the interface.</p> <p data-bbox="837 520 1403 583">You can configure a dynamic port profile based on the following parameters:</p> <ul data-bbox="837 617 1179 974" style="list-style-type: none"> <li>• LLDP System Name</li> <li>• LLDP Description</li> <li>• LLDP Chassis ID</li> <li>• Radius Username</li> <li>• Radius Filter-ID</li> <li>• MAC (Ethernet mac-address)</li> </ul> <p data-bbox="837 1010 1370 1073">In this example, the rule applies a port profile to all devices with the specified LLDP system name.</p> <div data-bbox="837 1108 1409 1766" style="border: 1px solid #ccc; padding: 10px;"> <p data-bbox="857 1125 1154 1146"><b>DYNAMIC PORT CONFIGURATION</b></p> <p data-bbox="857 1178 1370 1251">Apply port profiles to ports based on properties of connected clients. First matching rule will be applied. Port range must have dynamic configuration enabled.</p> <div data-bbox="857 1266 1409 1745" style="border: 1px solid #00aaff; padding: 5px;"> <p data-bbox="1084 1283 1390 1304" style="text-align: right;">New Rule <span style="float: right;">✓ ×</span></p> <p data-bbox="862 1346 1149 1377">Check <span style="border: 1px solid #ccc; padding: 2px;">LLDP System Name</span> <span style="float: right;">▼</span></p> <p data-bbox="862 1409 1370 1440"><input type="checkbox"/> Select the <span style="border: 1px solid #ccc; padding: 2px;">1st</span> <span style="float: right;">▼</span> segment (separated by <span style="border: 1px solid #ccc; padding: 2px;"> </span>)</p> <p data-bbox="862 1472 1338 1503"><input type="checkbox"/> Start at character offset <span style="border: 1px solid #ccc; padding: 2px;">0</span> (0 = first character)</p> <p data-bbox="862 1535 1008 1556">If text starts with</p> <p data-bbox="862 1566 1230 1598" style="border: 1px solid #ccc; padding: 2px;">D4:20:B0</p> <p data-bbox="862 1608 1062 1629">comma-separated values</p> <p data-bbox="862 1650 1094 1671">Apply Configuration Profile</p> <p data-bbox="862 1682 1390 1713" style="border: 1px solid #ccc; padding: 2px;">AP <span style="float: right;">default(1), trunk, edge ▼</span></p> </div> </div>

Table 5: Shared Elements Configuration Options (Continued)

Option	Notes
VRF	<p>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 <b>Enable Dynamic Configuration</b> check box on the Port Config tab in the Select Switches section of the switch template or in the Port Configuration section of the switch details page.</p> <p>It takes a couple of minutes for a port profile to be applied a port after a client is recognized, and a couple of minutes after that for the port profile assignment status to appear on the Mist portal.</p> <p>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).</p> <p>Dynamic port configuration on a switch is meant for establishing connection to IoT devices, APs, and user port endpoints. You should not use it for creating connection between switches, switches and routers, and switches and firewalls. Also, you should not enable Dynamic Port Configuration on the uplink port.</p> <p>With VRF, you can divide an EX Series switch into multiple virtual routing instances, effectively isolating the traffic within the network. You can define a name for the VRF, specify the networks associated with it, and include any additional routes needed. You can specify IPv4 or IPv6 addresses for the additional route.</p> <p><b>NOTE:</b> You can't assign the default network (VLAN ID = 1) to VRF.</p>

## Shared Elements—Port Profiles

In the "Shared Elements" on page 56 section, you can configure port profiles. These options appear when you click **Add Profile** or when you click a profile to edit.



### NOTE:

- For general information about profiles, see "[Port Profiles Overview](#)" on page 11.
- If you're working at the site level, you might see asterisks (\*) next to the port profile names. These port profiles were created in the switch template. If you click them, you'll see the settings in view-only mode. To make site-specific changes (affecting only this site and not the switch template itself), select **Override Template Defined Profile** and then edit the settings.

**Table 6: Port Profile Configuration Options**

Option	Notes
Name, Port Enabled, and Description	Basic settings to identify and enable the port.
Mode	<ul style="list-style-type: none"> <li>• Trunk—Trunk interfaces typically connect to other switches, APs, and routers on the LAN. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Specify the Port Network, VoIP Network (if applicable), and Trunk Networks.</li> <li>• Access—Default mode. Access interfaces typically connect to network devices, such as PCs, printers, IP phones, and IP cameras. In this mode, the interface can be in a single VLAN only. Specify the Port Network and the VoIP Network (if applicable).</li> <li>• Port Network and VoIP Network—Select the VLANs for this port.</li> </ul>

Table 6: Port Profile Configuration Options (Continued)

Option	Notes
Use dot1x authentication	<p>Select this option to enable IEEE 802.1X authentication for Port-Based Network Access Control. 802.1X authentication is supported on interfaces that are members of private VLANs (PVLANS).</p> <p>The following options are available if you enable dot1x authentication on a port:</p> <ul style="list-style-type: none"> <li>• <b>Allow Multiple Supplicants</b>—Select this option to allow multiple end devices to connect to the port. Each device is authenticated individually.</li> <li>• <b>Dynamic VLAN</b>—Specify dynamic VLANs that will be returned by the RADIUS server attribute 'tunnel-private-group-ID' or 'Egress-VLAN-Name'. This configuration enables a port to perform dynamic VLAN assignment.</li> <li>• <b>MAC authentication</b>—Select this option to enable MAC authentication for the port. When this option is selected, you can also specify an <b>Authentication Protocol</b>. If you specify a protocol, it must be used by supplicants to provide authentication credentials.</li> <li>• <b>Use Guest Network</b>—Select this option to use a guest network for authentication. Then select a <b>Guest Network</b> from the drop-down list.</li> <li>• <b>Bypass authentication when server is down</b>—If you select this option, clients can join the network without authentication if the server is down.</li> <li>• <b>Reauthentication interval</b>—In a switch port profile that uses dot1x authentication, you can configure a timer that controls how often a client reauthenticates itself with the RADIUS server. The recommended value is 6 to 12 hours (21600 to 43200 seconds). The default value is 65000 seconds.</li> </ul> <p>You need to also do the following for dot1x authentication to work:</p> <ul style="list-style-type: none"> <li>• Configure a RADIUS server for dot1x authentication from the Authentication Servers tile in the All Switches Configuration section of the template.</li> <li>• 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.</li> </ul> <p>Mouse-over the port to see the RADIUS-assigned VLAN field. Ports with dot1x enabled are assigned a new VLAN by the RADIUS server when the 802.1x authentication is successful. This view is especially useful when checking whether a given VLAN on a port has changed following dot1x authentication.</p>



Table 6: Port Profile Configuration Options (Continued)

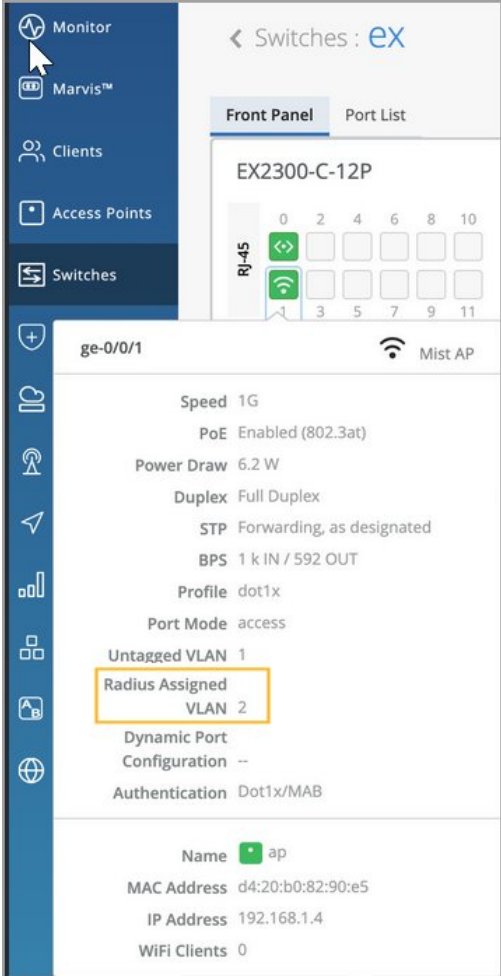
Option	Notes
	<p data-bbox="513 344 1089 373"><b>Figure 2: Radius Assigned VLAN on a Dot1x Port</b></p>  <p>The screenshot shows the configuration page for a Mist AP. The left sidebar contains navigation options: Monitor, Marvis™, Clients, Access Points, and Switches. The main content area shows the switch 'EX2300-C-12P' and a port list. The selected port is 'ge-0/0/1'. The configuration details for this port are as follows:</p> <ul style="list-style-type: none"> <li>Speed: 1G</li> <li>PoE: Enabled (802.3at)</li> <li>Power Draw: 6.2 W</li> <li>Duplex: Full Duplex</li> <li>STP: Forwarding, as designated</li> <li>BPS: 1 k IN / 592 OUT</li> <li>Profile: dot1x</li> <li>Port Mode: access</li> <li>Untagged VLAN: 1</li> <li>Radius Assigned VLAN: 2 (highlighted in yellow)</li> <li>Dynamic Port Configuration: --</li> <li>Authentication: Dot1x/MAB</li> </ul> <p>Additional information at the bottom of the configuration page includes:</p> <ul style="list-style-type: none"> <li>Name: ap</li> <li>MAC Address: d4:20:b0:82:90:e5</li> <li>IP Address: 192.168.1.4</li> <li>WiFi Clients: 0</li> </ul>
Speed	Keep the default setting, Auto, or select a speed
Duplex	Keep the default setting, Auto, or select a Half or Full.

Table 6: Port Profile Configuration Options (Continued)

Option	Notes
MAC Limit	<p>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.</p> <p>The default value: 0</p> <p>Supported range: 0 through 16383</p>
PoE	<p>Enable the port to support power over Ethernet (PoE).</p>
RSTP Edge	<p>Configure the port as a Rapid Spanning Tree Protocol (RSTP) edge port, if you want to enable Bridge Protocol Data Unit (BPDU) guard on a port. RSTP is enabled on ports to which clients that do not participate in RSTP are connected. This setting ensures that the port is treated as an edge port and guards against the reception of BPDUs. If you plug a non-edge device into a port configured with RSTP Edge, the port is disabled. In addition, the Switch Insights page generates a Port BPDU Blocked event. The Front Panel on the <a href="#">"Switch Details" on page 149</a> will also display a BPDU Error for this port.</p> <p>You can clear the port of the BPDU error by selecting the port on the Front Panel and then clicking <b>Clear BPDU Errors</b>.</p> <p>You should not enable RSTP Edge on the Uplink port.</p> <p>You can also configure RSTP Edge at the switch level, from the Port Profile section on the switch details page.</p>
RSTP Point-to-Point	<p>This configuration changes the interface mode to point-to-point. Point-to-point links are dedicated links between two network nodes, or switches, that connect one port to another.</p>
RSTP No Root Port	<p>This configuration prevents the interface from becoming a root port.</p>

Table 6: Port Profile Configuration Options (*Continued*)

Option	Notes
QoS	<p>Enable Quality of Service (QoS) for the port to prioritize latency-sensitive traffic, such as voice, over other traffic on a port.</p> <p><b>NOTE:</b> 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.</p> <p>You have the option to override the QoS configuration on the WLAN settings page (<b>Site &gt; WLANs &gt; <i>WLAN name</i></b>). To override the QoS configuration, select the <b>Override QoS</b> check box and choose a wireless access class. The downstream traffic (AP &gt; client) gets marked with the override access class value specified. The override configuration doesn't support upstream traffic (client &gt; AP).</p> <p>See also: "<a href="#">QoS Configuration</a>" on page 90.</p>
Storm Control	<p>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.</p> <p>For more information, see <a href="#">Understanding Storm Control</a>.</p>

Table 6: Port Profile Configuration Options (Continued)


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

Table 6: Port Profile Configuration Options (Continued)

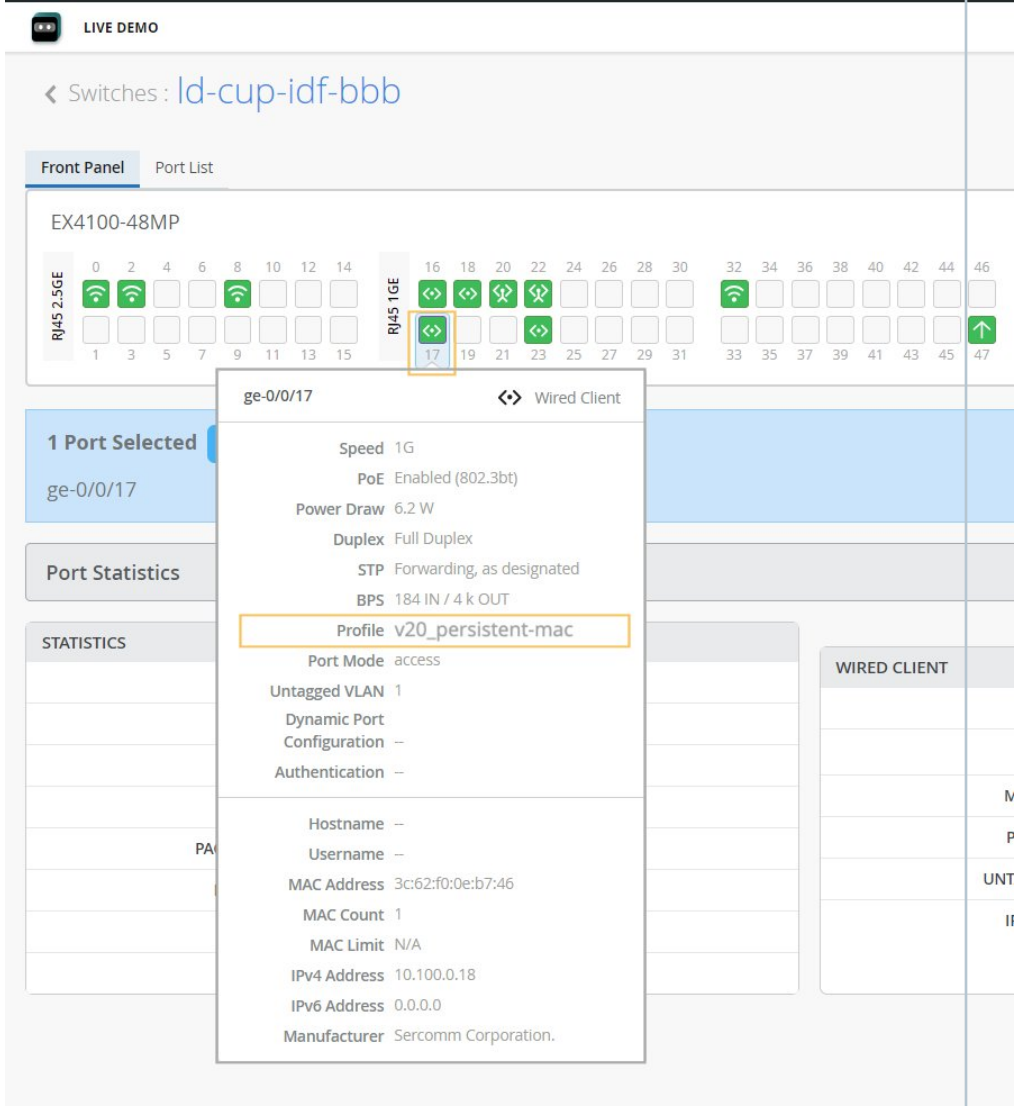
Option	Notes
	<p data-bbox="472 405 980 436"><b>Figure 3: Port Details Showing Sticky MAC</b></p>  <p data-bbox="472 1619 1421 1787">The configured MAC limit and number of MACs learned will appear after a few minutes, as dynamic learning on the interface progresses. In the Mist dashboard, only the maximum MAC address count is shown. However, you can see every MAC address a given interface has learned by opening a <b>Remote Shell</b> to the switch and running the following Junos CLI commands:</p>

Table 6: Port Profile Configuration Options (Continued)

Option	Notes
	<p>show ethernet-switching table persistent-learning show ethernet-switching table persistent-learning interface</p> <p>MAC count is a persistent value that remains until the MAC address is cleared (or until it is disabled in the Port Profile and then that configuration is pushed to the switch).</p> <p>To clear the MAC addresses on a given interface from the Mist dashboard, you need to be logged as Network Administrator or Super User. Then just select the port you want from the switch front panel (as shown in Figure 1) and click the <b>Clear MAC [Dynamic/Persistent]</b> button that appears.</p> <p>On the Switch Insights page, the event shows up as a <i>MAC Limit Reset</i> event.</p> <p>For more information on the front panel, see <a href="#">"Switch Details" on page 149</a>.</p>

## Select Switches Configuration

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

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

The screenshot shows the configuration interface for a switch rule. The tabs at the top are: Info, Port Config, IP Config, IP Config (OOB), Port Mirroring, and CLI Config. The 'Info' tab is selected. The form contains the following fields and options:

- Name:** ex-role2
- Applies to switch name
- Offset:** 0
- Applies to switch role
- Applies to switch model
- Role-2** (dropdown menu)
- EX2300-C-12P** (dropdown menu)
- letters, numbers, \_ or -

The various tabs are described in separate tables below.

**Table 7: Select Switches—Info Tab**

Option	Notes
Name	Enter a name to identify this rule.
Applies to switch name	Enable this option if you want this rule to apply to all switches that match the specified name. Then enter the text and the number of offset characters. For example, if you enter <i>abc</i> with an offset of 0, the rule applies to switches whose names start with <i>abc</i> . If the offset is 5, the rule ignores the first 5 characters of the switch name.
Applies to switch role	Enable this option if you want this rule to apply to all switches that have the same role. Enter the role by using lowercase letters, numbers, underscores ( <code>_</code> ), or dashes ( <code>-</code> ).
Applies to switch model	Enable this option if you want this rule to apply to all switches that have the same model. Then select the model.

Table 8: Select Switches—Port Config Tab

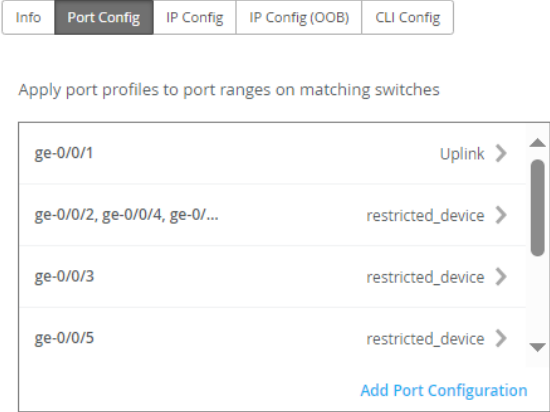
Option	Notes
Configuration List	<p>Click <b>Add Port Configuration</b>, or select a port configuration to edit.</p> 
Port IDs and Configuration Profile	Enter the port(s) to configure and the configuration profile to apply to them.



Table 8: Select Switches—Port Config Tab (Continued)

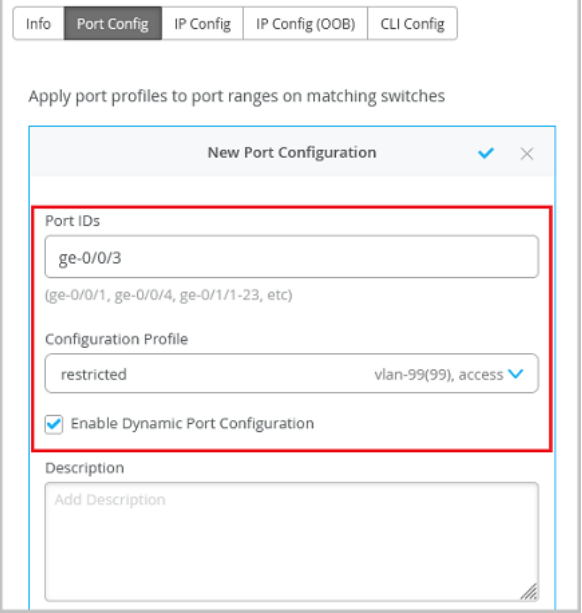
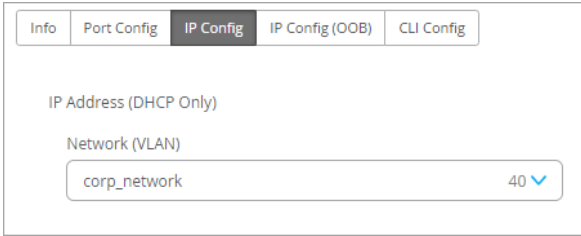
Option	Notes
Enable Dynamic Configuration	<p>When you enable this feature, a port profile is assigned based on defined attributes of the connected device. If the device matches the attributes, Mist assigns a matching dynamic profile to the device. But if the device doesn't match the attributes, it is placed in a specified VLAN.</p> <p>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.</p> 
Up/Down Port Alerts	<p>When you enable this feature, Juniper Mist monitors transitions between up and down states on these ports. If you enable this feature, also enable Critical Switch Port Up/Down on the Monitor &gt; Alerts &gt; Alerts Configuration page.</p>

Table 8: Select Switches—Port Config Tab (Continued)

Option	Notes
Port Aggregation	<p>When you enable this feature, Ethernet interfaces are grouped to form a single link layer interface. This interface is also known as a link aggregation group (LAG) or bundle.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Allow switch port operator to modify port profile	<p>When you enable this feature, users with the Switch Port Operator admin role can view and manage this configuration.</p>

Table 9: Select Switches Configuration—IP Config Tab

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

### Select Switches—IP Config (OOB) Tab

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

### Select Switches—Port Mirroring Tab

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

- **Input**—The source (an interface or network) of the traffic to be monitored. Along with the input, you can specify whether you want Mist to monitor the ingress traffic or the egress traffic for an interface.

If you want both ingress and egress traffic to be monitored, add two input entries for the same interface - one with the ingress flag and the other with the egress flag.

- **Output**—The destination interface to which you want to mirror the traffic. You cannot specify the same interface or network in both the input and output fields.

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

### Select Switches—CLI Config Tab

Enter additional CLI commands, as needed.

## Switch Policy Labels (Beta)

In this section, add GBP tags to identify groups of users and resources to reference in your switch policies.



**NOTE:** 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 example shows the tags that a user created, along with the policies that reference them.

The screenshot displays the Juniper Mist management interface. On the left is a navigation sidebar with icons for Monitor, Marvis™, Clients, Access Points, Switches, WAN Edges, Mist Edges, Location, Analytics, Site, and Organization. The main content area is divided into two sections:

- Switch Policy Labels (BETA):** This section contains a search bar and a table of 9 GBP Tags. The table has columns for NAME, TYPE, FROM, VALUE, and GBP TAG.
 

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
Desktop1	Static	Subnets	10.99.99.99	99
Developers	Static	Network	vlan1088	1088
Desktop2	Static	Subnets	10.99.99.99	99
- Switch Policy (BETA):** This section contains a search bar and a table of 4 Switch Policies. The table has columns for NO., NAME, USER/GROUP, and RESOURCE.
 

NO.	NAME	USER/GROUP	RESOURCE
1	Desktops	Desktop2	Desktop3, Desktop1
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

**NOTE:** This feature is currently available only to beta participants.

To get started, click **Add GBP tag**. Then enter a name, select the type, and enter the value. Then click Add at the lower-right corner of the screen.

When you enable a tag, you'll see on-screen alert about the impact on standalone switches and virtual chassis. Read the on-screen information before proceeding.

**NOTE:** If you configure 802.1X authentication with multiple-supplciant mode, the GBP tagging is MAC-based. If you configure 802.1X authentication with single-supplciant mode, the GBP tagging is port-based.

**Table 10: Switch Policy Label (GBP Tag) Configuration Options**

Option	Notes
Dynamic or Static	By default, Juniper Mist chooses the Dynamic option. If you select Static, specify a GBP tag source. It can be a MAC address, network, or an IP subnet.

Table 10: Switch Policy Label (GBP Tag) Configuration Options (*Continued*)

Option	Notes
GBP Tag	Enter value or GBP source tag for host-originated packets (range: 1 through 65535).

## Switch Policy (Beta)

Configure Group Based Policies (GBPs) that you can use in your campus fabric IP Clos deployments. When you create a policy, you use organization-level and site-level labels and your GBP tags (from the Switch Policy Labels section) to identify users who can or cannot access specified resources.



**NOTE:** 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 policy, using tags that were defined in the Switch Policy Labels section.

**Switch Policy Labels** BETA

GROUP BASED POLICY TAGS ⓘ

9 GBP Tags Add GBP Tag

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
Desktop1	Static	Subnets	10.99.99.99	99
Developers	Static	Network	vlan1088	1088
Desktop2	Static	Subnets	10.99.99.99	99

**Switch Policy** BETA

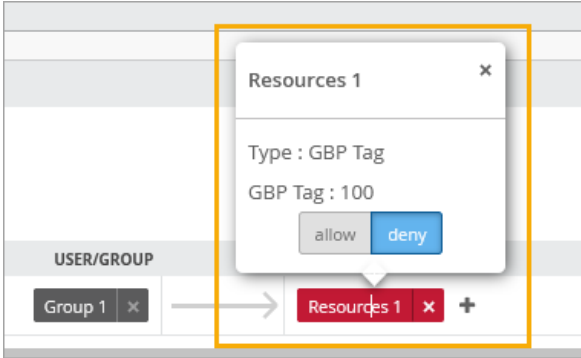
SWITCH POLICY

4 Switch Policies Add Switch Policy

NO.	NAME	USER/GROUP	RESOURCE
1	Desktops	Desktop2	Desktop3, Desktop1
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 get started, click **Add Switch Policy**. Then enter the settings, as described in the following table.

Table 11: Switch Policy Configuration Options

Option	Notes
USER/GROUP	Click + 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.
RESOURCE	<p>Click + 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.</p> <p>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 <b>deny</b>.</p> 

## Protection of Routing Engine

### IN THIS SECTION

- [Configure Protection of Routing Engine | 77](#)
- [Verify Protection of Routing Engine Configuration | 81](#)

The Protection of Routing Engine feature ensures that the Routing Engine accepts traffic only from trusted systems. Enabling this feature results in creation of a stateless firewall filter that discards all

traffic destined for the Routing Engine, except those from the specified trusted sources. Protecting the Routing Engine involves filtering incoming traffic on the router's lo0 interface. Enabling this feature on Juniper Switches is suggested as a best practice.

## Configure Protection of Routing Engine

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

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

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

**PROTECTION OF ROUTING ENGINE** ⓘ

Override Site/Template Settings

Protect Routing Engine  
 Enabled    Disabled

**Trusted Networks** ⓘ  
  
(Comma-separated list of IP addresses or CIDRs)

**Trusted Services** ⓘ  
 icmp  
 ssh

**Trusted IP/Protocol/Port** ⓘ Add IP/Protocol/Port

IP Address	Protocol	Port Range

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



**Add Trusted IP/Protocol/Port** [X]

IP Address  
[Text Input] /

(Comma-separated list of IP addresses or CIDRs)

Protocol  
[Dropdown: Any] ▾

Port Range  
[Text Input]

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

[Add] [Cancel]

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

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

To configure Protection of Routing Engine at the switch level:

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

**PROTECTION OF ROUTING ENGINE** ⓘ

Override Site/Template Settings

Protect Routing Engine

Enabled    Disabled

Trusted Networks ⓘ

(Comma-separated list of IP addresses or CIDRs)

Trusted Services ⓘ

icmp

ssh

Trusted IP/Protocol/Port ⓘ Add IP/Protocol/Port

IP Address	Protocol	Port Range

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

5. To add additional IP or IP Subnet that you want the switch to communicate with, enter the IP addresses in a comma separated format in the **Trusted Networks** field.
6. If you want the switch to respond to the SSH and ICMP services, select the **ssh** and **icmp** check boxes.

7. If you want the switch to respond to custom services (which are a specific combination of IP, Port and Protocol), follow the below steps:
  - a. Click **Add IP/Protocol/Port**.  
The Add Trusted IP/Protocol/Port window is displayed.
  - b. In the Add Trusted IP/Protocol/Port window, specify the IP Address, a Protocol, and an applicable Port Range.
  - c. Click **Add**.
8. Save the configuration.

### Configuration Commands (CLIs)

```
"set groups top firewall family inet filter protect_re term allow_mist_obssh from source-port
[ 2200 ]",
"set groups top firewall family inet filter protect_re term allow_mist_obssh then accept",
"set groups top firewall family inet filter protect_re term allow_dhcp from source-port [ 67
68 ]",
"set groups top firewall family inet filter protect_re term allow_dhcp from destination-port
[ 67 68 ]",
"set groups top firewall family inet filter protect_re term allow_dhcp from protocol udp",
"set groups top firewall family inet filter protect_re term allow_dhcp then accept",
"set groups top firewall family inet filter protect_re term allow_bgp from source-prefix-list
bgp_neighbors",
"set groups top firewall family inet filter protect_re term allow_bgp from source-prefix-list
bgp_vrf_neighbors",
"set groups top firewall family inet filter protect_re term allow_bgp from destination-port
[ 179 ]",
"set groups top firewall family inet filter protect_re term allow_bgp from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_bgp then accept",
"set groups top firewall family inet filter protect_re term allow_bfd from source-prefix-list
bgp_neighbors",
"set groups top firewall family inet filter protect_re term allow_bfd from source-prefix-list
bgp_vrf_neighbors",
"set groups top firewall family inet filter protect_re term allow_bfd from destination-port
[ 3784 4784 ]",
"set groups top firewall family inet filter protect_re term allow_bfd from protocol udp",
"set groups top firewall family inet filter protect_re term allow_bfd then accept",
"set groups top firewall family inet filter protect_re term allow_ntp from source-prefix-list
ntp_servers",
"set groups top firewall family inet filter protect_re term allow_ntp from destination-port
[ 123 ]",
"set groups top firewall family inet filter protect_re term allow_ntp from protocol udp",
```

```

"set groups top firewall family inet filter protect_re term allow_ntp then accept",
"set groups top firewall family inet filter protect_re term allow_dns from source-port [ 53 ]",
"set groups top firewall family inet filter protect_re term allow_dns from protocol [ tcp
udp ]", "set groups top firewall family inet filter protect_re term allow_dns then accept",
"set groups top firewall family inet filter protect_re term allow_radius from source-prefix-list
radius_servers",
"set groups top firewall family inet filter protect_re term allow_radius from destination-port
[ 1812 1813 ]",
"set groups top firewall family inet filter protect_re term allow_radius from protocol udp",
"set groups top firewall family inet filter protect_re term allow_radius then accept",
"set groups top firewall family inet filter protect_re term allow_tacacs from source-prefix-list
tacacs_servers",
"set groups top firewall family inet filter protect_re term allow_tacacs from destination-port
[ 49 ]",
"set groups top firewall family inet filter protect_re term allow_tacacs from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_tacacs then accept",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from source-
prefix-list snmp_clients",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from destination-
port [ 161 10161 ]",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from protocol
udp",
"set groups top firewall family inet filter protect_re term allow_snmp_clients then accept",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 10-216-192-1_32",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 100-100-100-2_32",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 8-8-8-8_32",
"set groups top firewall family inet filter protect_re term trusted_hosts then accept",
"set groups top firewall family inet filter protect_re term otherwise then discard",
"set groups top interfaces lo0 unit 0 family inet filter input protect_re",

```

## Verify Protection of Routing Engine Configuration

### IN THIS SECTION

- [Protection of Routing Engine \(Trusted Networks Configuration\) | 82](#)

- Protection of Routing Engine (Trusted Services Configuration) | 84

## Protection of Routing Engine (Trusted Networks Configuration)

### Configuration commands (CLI)

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

### APIs

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

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

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

```

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

```

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

```

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

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

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

{master:0}

```

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

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

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

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

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

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

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

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

## Protection of Routing Engine (Trusted Services Configuration)

### Configuration commands (CLI)

```
"set groups top interfaces lo0 unit 0 family inet filter input protect_re",
"set groups top firewall family inet filter protect_re term allow_mist_obssh from source-port
[ 2200 ]",
"set groups top firewall family inet filter protect_re term allow_mist_obssh then accept",
"set groups top firewall family inet filter protect_re term allow_dhcp from source-port [ 67
68 ]",
"set groups top firewall family inet filter protect_re term allow_dhcp from destination-port
[ 67 68 ]",
```

```
"set groups top firewall family inet filter protect_re term allow_dhcp from protocol udp",
"set groups top firewall family inet filter protect_re term allow_dhcp then accept",
"set groups top firewall family inet filter protect_re term allow_bgp from source-prefix-list
bgp_neighbors",
"set groups top firewall family inet filter protect_re term allow_bgp from source-prefix-list
bgp_vrf_neighbors",
"set groups top firewall family inet filter protect_re term allow_bgp from destination-port
[ 179 ]",
"set groups top firewall family inet filter protect_re term allow_bgp from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_bgp then accept",
"set groups top firewall family inet filter protect_re term allow_bfd from source-prefix-list
bgp_neighbors",
"set groups top firewall family inet filter protect_re term allow_bfd from source-prefix-list
bgp_vrf_neighbors",
"set groups top firewall family inet filter protect_re term allow_bfd from destination-port
[ 3784 4784 ]",
"set groups top firewall family inet filter protect_re term allow_bfd from protocol udp",
"set groups top firewall family inet filter protect_re term allow_bfd then accept",
"set groups top firewall family inet filter protect_re term allow_ntp from source-prefix-list
ntp_servers",
"set groups top firewall family inet filter protect_re term allow_ntp from destination-port
[ 123 ]",
"set groups top firewall family inet filter protect_re term allow_ntp from protocol udp",
"set groups top firewall family inet filter protect_re term allow_ntp then accept",
"set groups top firewall family inet filter protect_re term allow_dns from source-port [ 53 ]",
"set groups top firewall family inet filter protect_re term allow_dns from protocol [ tcp
udp ]", "set groups top firewall family inet filter protect_re term allow_dns then accept",
"set groups top firewall family inet filter protect_re term allow_radius from source-prefix-list
radius_servers",
"set groups top firewall family inet filter protect_re term allow_radius from destination-port
[ 1812 1813 ]",
"set groups top firewall family inet filter protect_re term allow_radius from protocol udp",
"set groups top firewall family inet filter protect_re term allow_radius then accept",
"set groups top firewall family inet filter protect_re term allow_tacacs from source-prefix-list
tacacs_servers",
"set groups top firewall family inet filter protect_re term allow_tacacs from destination-port
[ 49 ]",
"set groups top firewall family inet filter protect_re term allow_tacacs from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_tacacs then accept",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from source-
prefix-list snmp_clients",
"set groups top firewall family inet filter protect_re term allow_snmp_clients from destination-
port [ 161 10161 ]",
```



```

"set groups top firewall family inet filter protect_re term allow_snmp_clients from protocol
udp",
"set groups top firewall family inet filter protect_re term allow_snmp_clients then accept",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 10-216-192-1_32",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 100-100-100-2_32",
"set groups top firewall family inet filter protect_re term trusted_hosts from source-prefix-
list 8-8-8-8_32",
"set groups top firewall family inet filter protect_re term trusted_hosts then accept",
"set groups top firewall family inet filter protect_re term allow_ssh from destination-port
[ 22 ]",
"set groups top firewall family inet filter protect_re term allow_ssh from protocol tcp",
"set groups top firewall family inet filter protect_re term allow_ssh then accept",
"set groups top firewall family inet filter protect_re term allow_icmp from protocol icmp",
"set groups top firewall family inet filter protect_re term allow_icmp then accept",
"set groups top firewall family inet filter protect_re term otherwise then discard",

```

## APIs

```

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

```

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

```

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

```

```

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

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

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

```

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

```

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

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

```

```

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

```

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

```

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

```

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

## QoS Configuration

### IN THIS SECTION

- [Enable QoS on a Switch Port | 92](#)
- [Override QoS | 94](#)
- [Verify QoS Settings \(API\) | 94](#)
- [Verify QoS Configuration Through the CLI | 95](#)



**NOTE:** QoS configuration is part of the switch configuration workflow described in "[Configure Switches](#)" on [page 30](#). This topic provides more detailed information focused solely on QoS concept and configuration steps.

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

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

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

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

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

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

## Enable QoS on a Switch Port

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



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

To enable QoS on a switch port:

1. To configure QoS at the organization level, click **Organization** > **Switch Templates** > *template name*. Or, if you want to configure QoS at the switch level, click **Switches** > *switch name*.
2. From the Port Profile tile, select the port profile you want to update. Or if you want to create a new port profile, click **Add Profile**.
3. In the configuration, remember to select the **QoS** check box.

### New Port Profile

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

Name

Port Enabled

Enabled  Disabled

Description

---

Mode

Trunk  Access

Port Network (Untagged/Native VLAN)

 1 

---

VoIP Network

---

Use dot1x authentication

Speed

Duplex

Mac Limit

(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. You can do that from the Port Config tab in the Select Switches section of a switch configuration template (See ["Create a Switch Configuration Template" on page 31](#)) or from the Port Configuration section on the Switch Details page (["Switch Details" on page 149](#)).

## Override QoS

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

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

## Verify QoS Settings (API)

The following example has "enable\_qos": true set for the port profiles qos-test and uplink. This indicates that the port profile has QoS enabled.

```
"port_usages": {
  "qos-test": {
    "name": "qos-test",
    "mode": "access",
    "disabled": false,
    "port_network": "v110",
    "voip_network": null,
    "stp_edge": false,
    "all_networks": false,
    "networks": [],
    "port_auth": null,
    "speed": "auto",
    "duplex": "auto",
```

```

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

```

## Verify QoS Configuration Through the CLI

The following is a sample QoS configuration on a switch:

```

set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class bizapp-af2 loss-priority high code-points af22
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class bizapp-af2 loss-priority high code-points af23
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class bizapp-af2 loss-priority low code-points af21
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class bizapp-af3 loss-priority high code-points af32
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class bizapp-af3 loss-priority high code-points af33
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class bizapp-af3 loss-priority low code-points af31

```

```
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class bizapp-af3 loss-priority low code-points cs3
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class default-app loss-priority low code-points be
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class net-control loss-priority low code-points nc1
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class net-control loss-priority low code-points nc2
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class net-tools loss-priority high code-points af12
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class net-tools loss-priority high code-points af13
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class net-tools loss-priority low code-points af11
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class video loss-priority high code-points af42
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class video loss-priority high code-points af43
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class video loss-priority low code-points af41
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class video loss-priority low code-points cs4
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class voice loss-priority low code-points cs5
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default forwarding-
class voice loss-priority low code-points ef
set groups mist-qos-default class-of-service classifiers dscp dscp-classifier-default import
default
set groups mist-qos-default class-of-service forwarding-classes queue 0 default-app
set groups mist-qos-default class-of-service forwarding-classes queue 1 video
set groups mist-qos-default class-of-service forwarding-classes queue 2 bizapp-af3
set groups mist-qos-default class-of-service forwarding-classes queue 3 bizapp-af2
set groups mist-qos-default class-of-service forwarding-classes queue 4 net-tools
set groups mist-qos-default class-of-service forwarding-classes queue 5 voice
set groups mist-qos-default class-of-service forwarding-classes queue 7 net-control
set groups mist-qos-default class-of-service interfaces ge-0/0/0 scheduler-map sched-maps-default
set groups mist-qos-default class-of-service interfaces ge-0/0/0 unit 0 classifiers dscp dscp-
classifier-default
set groups mist-qos-default class-of-service interfaces ge-0/0/0 unit 0 rewrite-rules dscp dscp-
rewriter-default
set groups mist-qos-default class-of-service interfaces ge-0/0/9 scheduler-map sched-maps-default
set groups mist-qos-default class-of-service interfaces ge-0/0/9 unit 0 classifiers dscp dscp-
classifier-default
```

```
set groups mist-qos-default class-of-service interfaces ge-0/0/9 unit 0 rewrite-rules dscp dscp-rewriter-default
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewrite-default import default
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class bizapp-af2 loss-priority low code-point af21
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class bizapp-af3 loss-priority low code-point af31
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class default-app loss-priority low code-point be
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class net-control loss-priority low code-point nc1
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class net-tools loss-priority low code-point af11
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class video loss-priority low code-point af41
set groups mist-qos-default class-of-service rewrite-rules dscp dscp-rewriter-default forwarding-class voice loss-priority low code-point ef
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class bizapp-af2 scheduler bizapp-af2-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class bizapp-af3 scheduler bizapp-af3-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class default-app scheduler default-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class net-control scheduler net-control-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class net-tools scheduler net-tools-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class video scheduler video-scheduler
set groups mist-qos-default class-of-service scheduler-maps sched-maps-default forwarding-class voice scheduler voice-scheduler
set groups mist-qos-default class-of-service schedulers bizapp-af2-scheduler buffer-size percent 10
set groups mist-qos-default class-of-service schedulers bizapp-af2-scheduler priority low
set groups mist-qos-default class-of-service schedulers bizapp-af2-scheduler transmit-rate percent 10
set groups mist-qos-default class-of-service schedulers bizapp-af3-scheduler buffer-size percent 10
set groups mist-qos-default class-of-service schedulers bizapp-af3-scheduler priority low
set groups mist-qos-default class-of-service schedulers bizapp-af3-scheduler transmit-rate percent 10
set groups mist-qos-default class-of-service schedulers default-scheduler buffer-size remainder
```

```

set groups mist-qos-default class-of-service schedulers default-scheduler priority low
set groups mist-qos-default class-of-service schedulers default-scheduler transmit-rate remainder
set groups mist-qos-default class-of-service schedulers net-control-scheduler buffer-size
percent 3
set groups mist-qos-default class-of-service schedulers net-control-scheduler priority low
set groups mist-qos-default class-of-service schedulers net-control-scheduler transmit-rate
percent 3
set groups mist-qos-default class-of-service schedulers net-tools-scheduler buffer-size percent 3
set groups mist-qos-default class-of-service schedulers net-tools-scheduler priority low
set groups mist-qos-default class-of-service schedulers net-tools-scheduler transmit-rate
percent 3
set groups mist-qos-default class-of-service schedulers video-scheduler buffer-size percent 8
set groups mist-qos-default class-of-service schedulers video-scheduler priority low
set groups mist-qos-default class-of-service schedulers video-scheduler transmit-rate percent 8
set groups mist-qos-default class-of-service schedulers voice-scheduler buffer-size percent 10
set groups mist-qos-default class-of-service schedulers voice-scheduler priority strict-high
set groups mist-qos-default class-of-service schedulers voice-scheduler shaping-rate percent 10

```

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

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

```

root@ex2300-home> show interfaces ge-0/0/0 extensive
.....
Queue counters:      Queued packets  Transmitted packets  Dropped packets
0                   0                0                    0
1                   0                0                    0
2                   0                0                    0
3                   0                0                    0
4                   0                0                    0
5                   0                0                    0
7                   0                0                    0
Queue number:      Mapped forwarding classes
0                 default-app
1                 video
2                 bizapp-af3
3                 bizapp-af2
4                 net-tools
5                 voice
7                 net-control
.....

```

```

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

```

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

```

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

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

```

```

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

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

```

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

```

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

```

```

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

```

See also: [Example: Configuring CoS on EX Series Switches](#)

## Configure SNMP on Switches

### IN THIS SECTION

- [Configure SNMP at the Organization Level | 101](#)
- [Configure SNMP at the Site Level | 103](#)
- [Configure SNMP at the Device Level | 105](#)



**NOTE:** SNMP configuration is part of the switch configuration workflow described in "[Configure Switches](#)" on page 30. This topic provides more detailed information focused solely on SNMP configuration steps.

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

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

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

## Configure SNMP at the Organization Level

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

To configure SNMP at the organization level:

1. Click **Organization > Switch Templates**.
2. Click **Create Template** if you want to configure SNMP as part of a new switch template. For more information, see "[Create a Switch Configuration Template](#)" on page 31.  
Or, if you want to update an existing template with SNMP configuration, click that template to open it.
3. On the SNMP tile (in the All Switches Configuration section), click the **Enabled** check box.



**SNMP**

SNMP fields cannot be all empty

Enabled     Disabled

V2     V3

<b>General</b>	Clients	Trap Groups	Community	Views
----------------	---------	-------------	-----------	-------

Name

Location

Contact

Description

Source Address

None
▼

4. Specify the SNMP information, which includes General information, Clients, Trap Groups, Community, and Views. Mist supports the SNMP versions V2 and V3. For the field descriptions, refer to the SNMP information in the **All Switches Configuration Options** table in ["All Switches" on page 38](#).

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

5. Save the template.
6. Assign the template to a site. For instructions to do so, see ["Assign a Template to Sites" on page 32](#).

## Configure SNMP at the Site Level

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

To configure SNMP at the site level:

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

**SNMP**

SNMP fields cannot be all empty

Override Configuration Template

Enabled     Disabled

V2     V3

General
Clients
Trap Groups
Community
Views

Name

Location

Contact

Description

Source Address

None
▼

6. Specify the SNMP information, which includes General information, Clients, Trap Groups, Community, and Views. Mist supports the SNMP versions V2 and V3. For the field descriptions, refer to the SNMP information in the **All Switches Configuration Options** table in ["All Switches" on page 38](#).

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

7. Save the configuration.

## Configure SNMP at the Device Level

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

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



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

To configure SNMP at the switch level:

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

**SNMP**

SNMP fields cannot be all empty

Override Site/Template Settings

Enabled     Disabled

V2     V3

**General**
Clients
Trap Groups
Community
Views

Name

Location

Contact

Description

Source Address

None
▼

6. Specify the SNMP information, which includes General information, Clients, Trap Groups, Community, and Views. Mist supports the SNMP versions V2 and V3. For the field descriptions, refer to the SNMP information in the **All Switches Configuration Options** table in ["All Switches" on page 38](#).

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

7. Save the configuration.

# Configure DHCP Server or Relay on a Switch

## IN THIS SECTION

- Prerequisites | 107
- Configure DHCP Server | 107
- Configure DHCP Relay | 114

This chapter lists the steps that are required to configure DHCP server or relay on a switch.

## Prerequisites

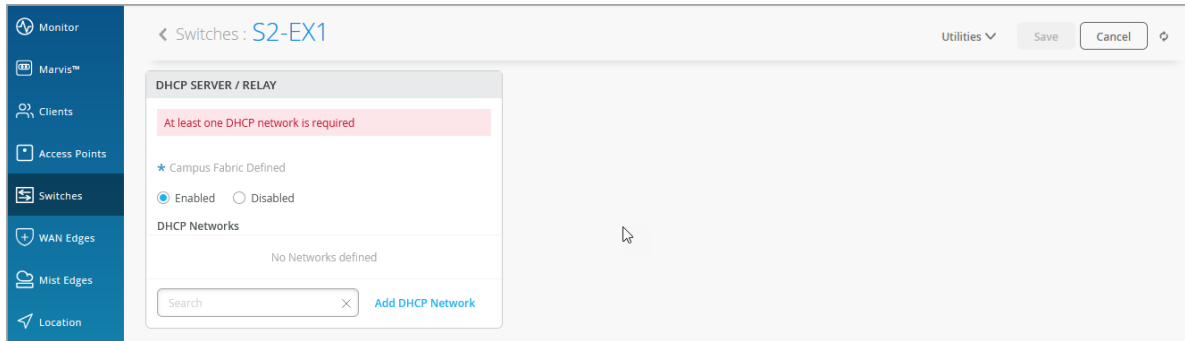
Before configuring the DHCP server or relay, ensure the following:

- The VLAN for which DHCP server will be configured on switch is assigned to the ports connecting to the DHCP clients. You can do this by applying a relevant port profile to the port. For more information about port profiles, see ["Port Profiles Overview" on page 11](#).
- The switch has a **Static** IP Configuration or Additional IP configuration for the network to run DHCP server. Or the L3 Interface is configured for DHCP server or relay.

## Configure DHCP Server

To configure DHCP server on a switch:

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



3. On the **DHCP Server / Relay** tile, select the **Enabled** check box and then click **Add DHCP Network**. The DHCP Server / Relay configuration window is displayed.
4. Ensure that **Server** is selected as configuration **Type**.
5. From the Network drop-down list, select a network for the DHCP server.
6. Specify the following fields:
  - IP Start—The starting IP address within the DHCP IP address assignment pool.
  - IP End—The ending IP address within the DHCP IP address assignment pool.
  - Gateway—Default gateway for the DHCP client. Usually, it is the switch IRB IP address for the corresponding network.
  - DNS Servers (Optional)—You can add up to three DNS IP addresses in a comma separated format.
  - DNS Suffix (Optional)—You can add up to three domain suffixes in a comma separated format.

DHCP SERVER / RELAY

Add DHCP Network
✓ ✕

Type

Server    Relay

Network ⓘ

vl10
10 ▼

IP Start

10.0.0.100

IP End

10.0.0.200

Gateway

10.0.0.10

DNS Servers

8.8.8.8,8.8.4.4,1.0.0.1

(Comma-separated IPs and Max 3)

DNS Suffix

test.net,tesing1.net,testing3.net

(Comma-separated domains and Max 3)

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

7. Save the changes.

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



```

"dhcpd config": {
  "enabled": true,
  "vl10": {
    "type": "server",
    "ip start": "10.0.0.100",
    "ip end": "10.0.0.200",
    "gateway": "10.0.0.10",
    "dns servers": [
      "8.8.8.8",
      "8.8.4.4",
      "1.0.0.1"
    ],
    "dns suffix": [
      "test.net",
      "tesing1.net",
      "testing3.net"
    ]
  }
},
}

```

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

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

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

The address assignment pool configuration looks like the below:

```

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

```

```
array string test.net",  
"set groups top access address-assignment pool vl10 family inet dhcp-attributes option 15  
array string tesing1.net",  
"set groups top access address-assignment pool vl10 family inet dhcp-attributes option 15  
array string testing3.net,"
```

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

```

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

```

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

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

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

```
Client IP Address: 10.0.0.100
Hardware Address: 5c:5b:35:f1:c9:00
State: BOUND(LOCAL_SERVER_STATE_BOUND)
Protocol-Used: DHCP
Lease Expires: 2022-06-16 18:10:16 IST
Lease Expires in: 85771 seconds
Lease Start: 2022-06-15 18:10:15 IST
Last Packet Received: 2022-06-15 18:10:16 IST
Incoming Client Interface: irb.10:ge-0/0/4.0
Server Identifier: 10.0.0.10
Session Id: 11
Client Pool Name: v110
Client IP Address: 10.0.0.102
Hardware Address: d4:20:b0:82:92:cf
State: BOUND(LOCAL_SERVER_STATE_BOUND)
Protocol-Used: DHCP
Lease Expires: 2022-06-16 18:11:11 IST
Lease Expires in: 85826 seconds
Lease Start: 2022-06-15 18:10:23 IST
Last Packet Received: 2022-06-15 18:11:11 IST
Incoming Client Interface: irb.10:ge-0/0/5.0
Server Identifier: 10.0.0.10
Session Id: 13
Client Pool Name: v110
```

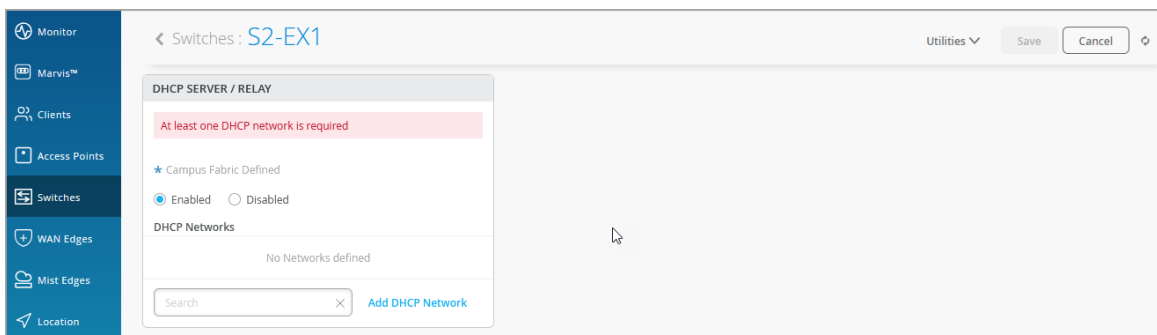
Here are some useful Junos CLI commands:

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

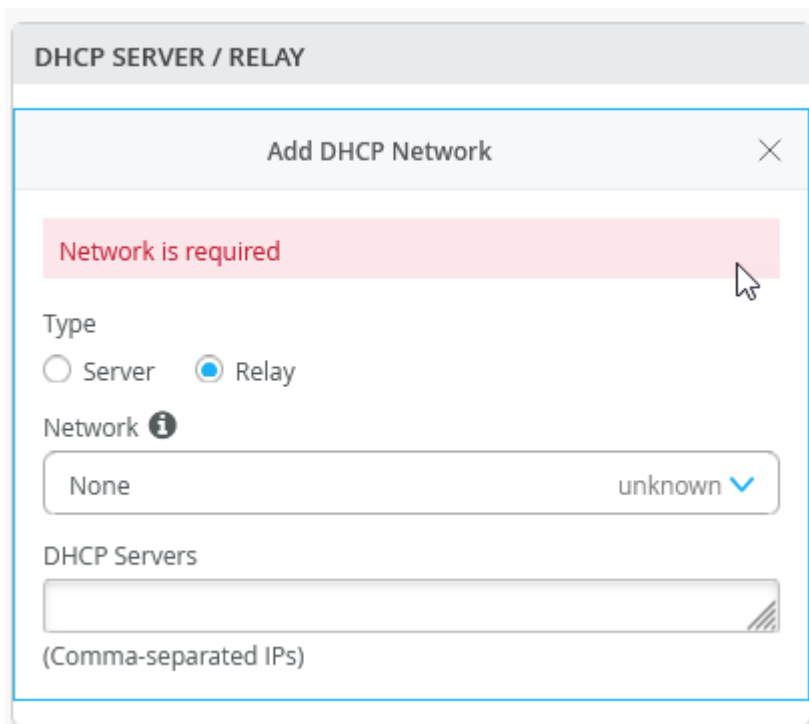
## Configure DHCP Relay

To configure DHCP relay on a switch:

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



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



5. From the Network drop-down list, select a network for the DHCP relay.
6. In the DHCP Servers field, configure the IP address for the remote DHCP server. You can add up to three IP addresses in a comma separated format.

7. Save the changes.

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

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

```
"dhcpd_config": {
  "enabled": true,

  "v111": {
    "type": "relay",
    "servers": [
      "192.168.1.1"
    ]
  }
}
```

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

The configuration looks like the below:

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

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

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

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

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

10. Verify DHCP relay binding on the switch.

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

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

Here are some useful Junos CLI commands:

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

See also: [DHCP Relay Agent](#).

## OSPF Configuration for Switches

### IN THIS SECTION

- [Example: Configure Basic OSPF in two EX Devices | 118](#)
- [CLI Commands | 122](#)



- [OSPF Events | 124](#)
- [Configure OSPF at the Organization Level | 124](#)
- [Configure OSPF at the Site Level | 125](#)

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

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

## Example: Configure Basic OSPF in two EX Devices

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

1. Navigate to the switch by clicking **Switches** > *Switch Name*.  
The switch details page appears.
2. Configure network by following the steps below
  - a. Navigate to the NETWORKS tile, and click Add Network.  
The New Network window appears.
  - b. Enter a network name (for example **vlan20**), VLAN ID, and subnet.
  - c. Click the check mark at the upper right of the New Network window to save the configuration.

The screenshot shows a web interface for configuring networks. At the top, there's a 'NETWORKS' section with a sub-header 'Named VLAN IDs that can be used by Port Profiles' and a note '\* Site, Template, or System defined'. Below this is a modal window titled 'Edit Network' with a checkmark and a close button. Inside the modal, there are two input fields: 'Name' with the value 'vlan20' and 'VLAN ID' with the value '20'. Below the 'VLAN ID' field, there is a small text '(1 - 4094 or {{siteVar}})'. The 'Name' field is highlighted with a red box, and the 'VLAN ID' field is also highlighted with a red box.

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

- a. On the PORT PROFILES window, click **Add Profile**.  
The New Port Profile window appears.
- b. Enter the profile details. In this example, name this profile as **vlan20portprofile**. In the port profile, you must include the network (**vlan20**) that you created in the previous step.
- c. Click the check mark at the upper right of the New Port Profile window to save the configuration.

PORT PROFILES

Port configuration for a set of related ports  
★ Site, Template, or System defined

New Port Profile

Name  
vlan20portprofile

Port Enabled  
 Enabled  Disabled

Mode  
 Trunk  Access

Port Network (Untagged/Native VLAN)  
vlan20 20

VoIP Network  
None

4. Attach the port profile to the port that is connected to the other switch. To do that:
  - a. On the PORT CONFIGURATION tile, click **Add Port Configuration**.  
The New Port Configuration window appears.
  - b. Enter the configuration details. Remember to specify the relevant interface in the Port ID field and include the **vlan20portprofile** in the Configuration Profile field.
  - c. Click the check mark at the upper right of the New Port Configuration window to save the configuration.

PORT CONFIGURATION

Port Profile Assignment  
★ Site, Template, or System defined

New Port Range

Port Aggregation

Port IDs  
ge-0/0/5  
(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Configuration Profile  
vlan20portprofile vlan20 (20)

Enable Dynamic Configuration

Description  
Add Description

5. Add an IP address to the network which you created earlier (vlan20). To do that:
  - a. On the IP CONFIGURATION tile, click **Add IP Configuration**.

The New IP Configuration window appears.

- b. Select Static as address type.
  - c. Specify an IP address and subnet mask. In this case, let's use 20.1.1.1/24 on one switch and 20.1.1.2/24 on the other switch.
  - d. From the Network (VLAN) drop-down list, select the network (vlan20) that you configured earlier.
  - e. Click the check mark at the upper right of the New IP Configuration window to save the configuration.
6. On the OSPF tile, configure an OSPF area by following the steps below:

- a. Click **Add Area**.

The New Area window appears.

- b. Specify the following details:
    - Area—Enter the area number (Range: 0 to 255).
    - Type—Specify an area type from the drop-down list.
    - Include Loopback—Select this check box if you want to include the loopback interface in the OSPF area.
  - c. Click **Add OSPF Network**.
- The Add OSPF Network window appears.

The screenshot shows the 'Add OSPF Network' configuration window. It includes the following settings:

- Network (VLAN): default
- Interface Type: broadcast
- Authentication Type: none
- Metric: 1-65535
- BFD Interval: 1-255000 milliseconds
- Enable Timers:
- Passive:
- OSPF Configuration:  Enabled,  Disabled

- d. Specify the network and the additional details as listed below:
- Network (VLAN)—From the drop-down list, select the network (vlan20) that you defined earlier.
  - Interface Type—Select an interface type. The following options are available: broadcast, p2p, and p2mp.
  - Authentication Type—Choose an authentication type from the following options: md5, password, and none. If you choose none, ensure that you select that option in both the switches.
  - Key—(Applicable if the Authentication Type chosen is md5). Specify a key for md5 authentication. These should be the same in both the devices.
  - Value—(Applicable if the Authentication Type chosen is md5). Specify a value for the specified md5 key. These should be the same in both the devices.
  - Password—(Applicable if the Authentication Type chosen is password). The password should be the same in both the switches for the OSPF neighborship to be up.
  - Metric—Specify the cost of an OSPF interface.
  - BFD Interval—Specify the interval at which the device exchanges BFD packets with its peer. Range: 1 through 255000 (in milliseconds).
  - Enable Timers—This option allows you to configure Hello Interval and Dead Interval.
  - Hello Interval—(Applicable if Enable Times is selected) Specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. By default, the routing device sends hello packets every 10 seconds.

- **Dead Interval**—(Applicable if **Enable Times** is selected) Specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. By default, the routing device waits 40 seconds (four times the hello interval).
  - **Passive**—Select this check box to advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.
- e. Click the check mark on the upper right of the **Add OSPF Network** window to save the OSPF Network information.  
You are returned to the **New Area** window.
  - f. Click the check mark on the upper right of the **New Area** window to save the OSPF Area information.
7. Select the **Enabled** check box under the **OSPF Configuration** section on the **OSPF** tile.
  8. On the **Routing** tile, enter the router ID.
  9. Click **Save** on the upper right of the switch details page to save the configuration on the switch.
- Carry out the above steps on both the switches to establish the OSPF neighborhood between them.

## CLI Commands

When Interface Type is p2p and Authentication Type is md5:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 authentication md5 1 key $9$cXyK8ws...",
"set groups top protocols ospf area 0 interface irb.20 interface type p2p"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2p and Authentication Type is none:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
```

```
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 interface type p2p"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2p and Authentication Type is password:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 authentication simple-password
$9$cXyK8ws...",
"set groups top protocols ospf area 0 interface irb.20 interface type p2p"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2mp and Authentication Type is md5:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 authentication md5 1 key $9$cXyK8ws...",
"set groups top protocols ospf area 0 interface irb.20 interface type p2mp"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2mp and Authentication Type is none:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
```

```
"set groups top protocols ospf area 0 interface irb.20 interface type p2mp"
"set groups top routing-options router-id 11.1.1.1",
```

When Interface Type is p2mp and Authentication Type is password:

```
"set apply-groups top"
"set groups default interfaces <*> unit 0 family ethernet-switching vlan members [ default ]",
"set groups top forwarding-options storm-control-profiles default all",
"set groups top interfaces lo0 unit 0 family inet address 11.1.1.1/32",
"set groups top poe interface all",
"set groups top protocols ospf area 0 interface irb.20 authentication simple-password
$9$cXyK8ws...",
"set groups top protocols ospf area 0 interface irb.20 interface type p2mp"
"set groups top routing-options router-id 11.1.1.1",
```

## OSPF Events

The following image shows examples of OSPF events:

Switch Events		43 Total	39 Good	1 Neutral	3 Bad
OSPF Neighbor Up	01:26:50.000 PM, Jun 30				
OSPF Neighbor Down	01:26:49.000 PM, Jun 30				
OSPF Neighbor Up	01:26:31.000 PM, Jun 30				

## Configure OSPF at the Organization Level

To configure OSPF at the organization template level:

1. Click **Organization > Switch Templates**.
2. Open the organization template you want to modify.
3. Navigate to the OSPF tile and click **Add Area**.
4. Follow the steps listed in ["Example: Configure Basic OSPF in two EX Devices" on page 118](#).

## Configure OSPF at the Site Level

To configure OSPF at the site template level:

1. Click **Site** > **Switch Configuration**.
2. Open the site template you want to modify.
3. Navigate to the OSPF tile and click **Add Area**.
4. Follow the steps listed in "[Example: Configure Basic OSPF in two EX Devices](#)" on page 118.

# 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 | 125](#)
- [Update Switch Configuration Settings at the Site Level | 126](#)
- [Add or Delete a CLI Configuration | 127](#)

## 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 31.
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 31.

## 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.`

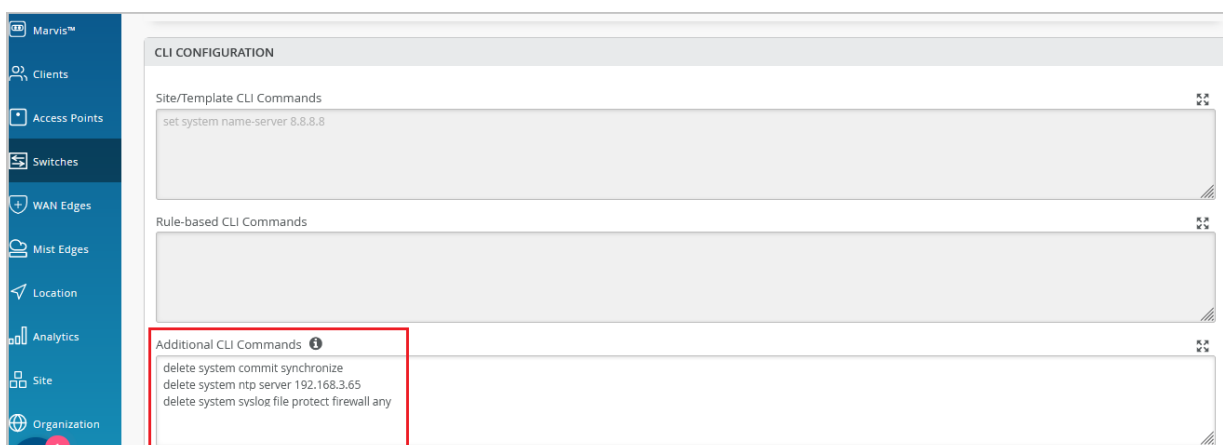
To add or delete a CLI configuration:

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.



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 149.
- 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

- [Free Up Storage Space on Your Switch | 129](#)
- [Upgrade the Junos OS Software on Your Switch | 131](#)

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.
- (Optional) A recovery snapshot stored on the OAM volume. See "[Switch Utilities](#)" on page 161 for details about the snapshot.

## 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

### Supported Devices

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).

### Available Versions

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

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

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

## Selecting a Release

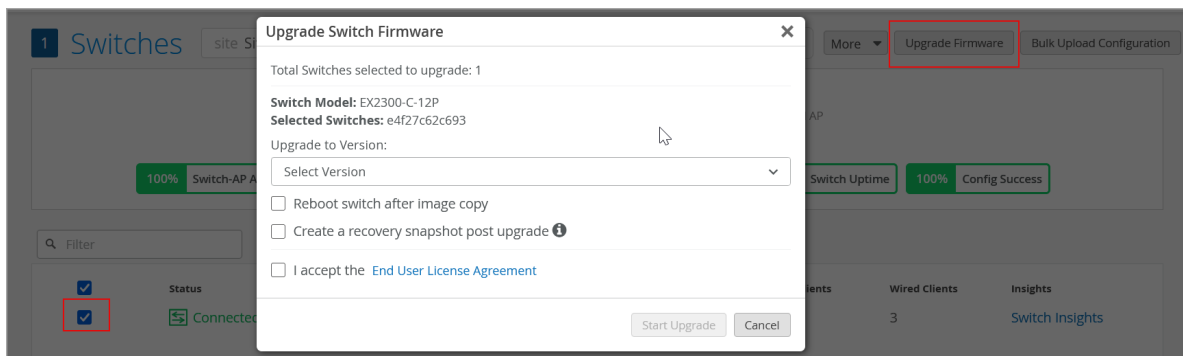
For more information about releases, consult these topics:

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

## Upgrading the Junos OS from Juniper Mist

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 Connectivity"](#) on page 261.
3. From the **List** tab, select the switch that requires a software upgrade, and then click **Upgrade Firmware**. You can select one or more switches for upgrade.



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 149).

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**. The drop-down list displays the suggested software version for the selected switch, along with all the applicable versions.

**Upgrade Switch Firmware** [X]

Total Switches selected to upgrade: 1

**Switch Model:** EX4100-F-12T  
**Selected Switches:** a4 10

Upgrade to Version:

Select Version [v]

**Suggested**

22.4R3.25

All

23.4R1.10

23.4R1.9

23.2R2.21

23.2R2.3

23.2R1-S2.5

23.2R1.13

22.4R3.25

22.4R2-S2.6

If you don't see the software version you are looking for, write to [support@mist.com](mailto: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**.

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

Select **I accept End User Agreement**.



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

The screenshot shows the Mist Switches interface for site 'Mist\_Office'. At the top, there are summary statistics: 2 Switches, 0 Mist APs, 2 Wired Clients, 0 Wireless Clients, and 0 W Total AP Power. Below these are filters for Switch-AP Affinity, PoE Compliance, VLANs, Version Compliance, and a green '100% Switch Uptime' indicator. The main table lists two switches:

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:fe: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.

The screenshot shows the Mist Switches interface with the 'Table Settings' dialog box open. The dialog box has 15 columns, each with a checkbox. The 'Status' checkbox is checked and highlighted with a red box. The background shows the same switch list as the previous screenshot, but the 'Status' column is currently hidden.

Table Settings
1. <input checked="" type="checkbox"/> Status
2. <input checked="" type="checkbox"/> Name
3. <input checked="" type="checkbox"/> IP Address
4. <input checked="" type="checkbox"/> Mist APs
5. <input checked="" type="checkbox"/> MAC Address
6. <input checked="" type="checkbox"/> Wired Clients
7. <input checked="" type="checkbox"/> Wireless Clients
8. <input checked="" type="checkbox"/> Model
9. <input checked="" type="checkbox"/> Version
10. <input checked="" type="checkbox"/> Total Power Draw
11. <input checked="" type="checkbox"/> Description
12. <input checked="" type="checkbox"/> Uptime
13. <input checked="" type="checkbox"/> Managed
14. <input checked="" type="checkbox"/> Role
15. <input checked="" type="checkbox"/> 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 149 page and click the **Switch Insights** link on the **Properties** tile.

The screenshot shows the 'Switch Events' section with 143 total events. The events are categorized as 131 Good, 0 Neutral, and 8 Bad. The 'All event Types' and 'All switch ports' filters are set to their default values. The following table shows the events:


Event Type	Device	Timestamp
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:



See also: ["Free Up Storage Space on Your Switch" on page 129](#)

- 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:

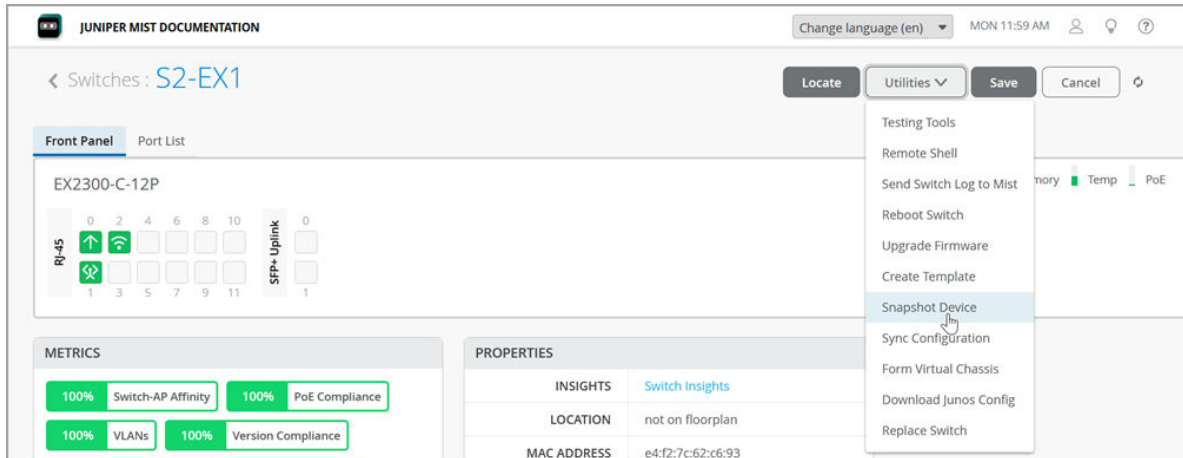
OC FWUPDATE WRITEFAILED. See also: [\[EX/QFX\] Certificate errors - Cannot validate Junos Image : Format error in certificate.](#)

## Create Recovery Snapshot for a Switch

You can create a recovery snapshot for a switch when needed. A recovery snapshot stored in OAM (Operations, Administration, and Maintenance) volume holds a full backup that can be used in case something goes wrong with Junos volume.

To create a recovery snapshot

1. Click **Switches** to go to the list of switches.
2. Locate the switch for which you want to create the snapshot and then click it to open the switch details page.
3. On the switch details page, click **Utilities > Snapshot Device**.



The **Snapshot Device** confirmation window appears.

4. On the **Snapshot Device** confirmation window, click **Request Snapshot**.

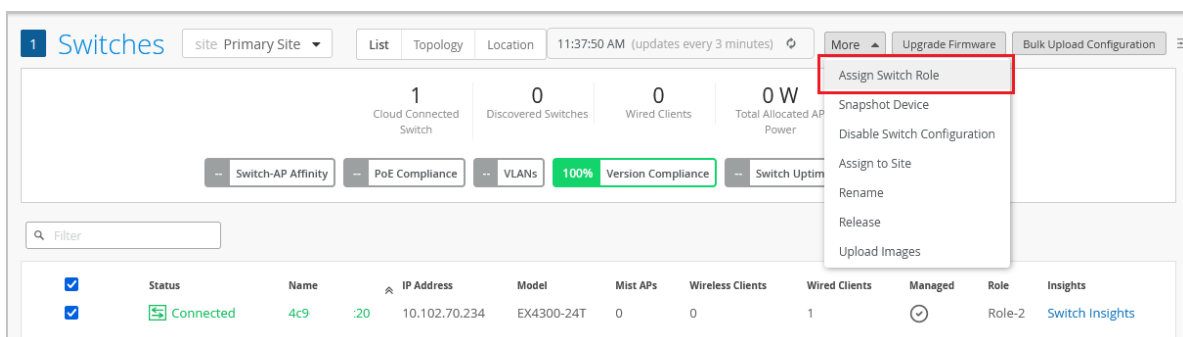
## Assign a Role to Switches

You can select and apply a role to an individual switch from the switch details page or from the switch list. This feature ensures that a switch does not inherit a role that does not exist. The role you assign to a switch should exist in the configuration template associated with the selected switch. A switch can have only one role at a time.

You can create new switch roles as part of Select Switches Configuration rules in a switch template (**Organization > Switch Templates** or **Sites > Switch Configuration**). See also: "[Configure Switches](#)" on page 30.

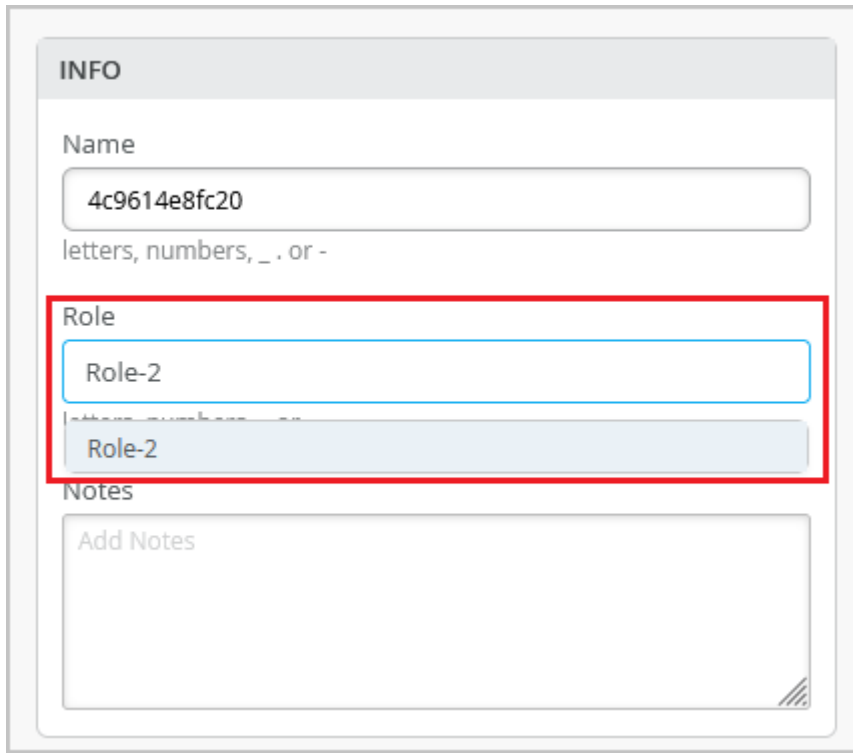
To assign a role to a switch:

1. Click **Switches** to go to the switch list.
2. Identify the switch that you want to assign a role to and then select the switch check box.
3. On the **More** menu, click **Assign Switch Role**.



4. In the **Role** field, specify the switch role and then click **OK**. Click inside the Role field to view all the available roles. The Role field displays the switch roles that exist in the configuration template associated with the selected switch.

Alternatively, navigate to the switch details page (**Switches > Switch-Name**), specify the switch role in the INFO section, and save the configuration (see the image below).



The image shows a configuration page for a switch, specifically the 'INFO' section. The 'Name' field contains the value '4c9614e8fc20'. Below it, the 'Role' field is highlighted with a red border. The 'Role' dropdown menu is open, showing 'Role-2' as the selected option. Below the 'Role' field is a 'Notes' section with a text area labeled 'Add Notes'.


## Locate a Switch by LED

To find a given switch or Virtual Chassis (VC) member on a rack full of switches, you can enable the **Locate** beacon to illuminate an LED on the switch. For VCs, this means you can **Locate** the primary, backup, or a given linecard member (only one member can be located at a time).

To turn on the LED beacon to identify the switch,

1. In the main menu, select **Switches** to open the switches inventory page.
2. Choose the switch you want to locate by clicking its name in the list of switches:
  - Click the **Locate** button at the top of the page to start the beacon, and again when you want to stop it.
  - To locate a member of VC, select the member in the drop down that appears.

Figure 4: The Mist LED on an EX Series Chassis



## 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 149 page appears.
3. Select the **Replace Switch** option from the **Utilities** drop-down list on the details.

The screenshot displays the Mist portal interface for a switch named 'ld-cup-idf-b-sw1'. The main content area shows a 'Front Panel' view of the EX3400-48P switch with a grid of ports. Below this, there are 'METRICS' and 'PROPERTIES' sections. The 'Utilities' dropdown menu is open, showing options like 'Testing Tools', 'Remote Shell', 'Send Switch Log to Mist', 'Reboot Switch', 'Upgrade Firmware', 'Create Template', 'Snapshot Device', 'Download junos Config', and 'Replace Switch'. The 'Replace Switch' option is highlighted with a red box.

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

<input type="checkbox"/> Role	<input type="checkbox"/> IP configuration (Out of Band)
<input type="checkbox"/> IP Configuration	<input type="checkbox"/> NTP
<input type="checkbox"/> Port Configuration	<input type="checkbox"/> Radius Configuration
<input type="checkbox"/> OSPF Areas	<input type="checkbox"/> DNS Servers
<input type="checkbox"/> DNS Suffix	<input type="checkbox"/> Static Route
<input type="checkbox"/> CLI Configuration	<input type="checkbox"/> DHCP Snooping
<input type="checkbox"/> Additional IP Configuration	<input type="checkbox"/> 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 30](#).

#### 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.

# Disable Remote Shell Access to Switches and Gateway Devices

Mist provides an option to turn off the remote shell access to switches and gateway devices in an organization. This setting is available at the organization level.

To turn off remote shell access:

1. Navigate to the Switch Management tile on the **Organization > Settings** page.
2. Select **Disable Remote Shell Access**.



## Connect a Switch to Mist Cloud via a Proxy Server Using Cloudx

### IN THIS SECTION

- [Connect a Switch to Mist Cloud via a Dynamic Proxy Server | 142](#)
- [Connect a Switch to Mist Cloud via a Static Proxy Server | 142](#)

This document explains the procedure to connect an EX Series switch to the Mist cloud via a proxy server (on-prem) directly without using a Mist Edge.



## Connect a Switch to Mist Cloud via a Dynamic Proxy Server

Before you connect the switch to the Mist cloud via a dynamic proxy server, ensure that the following prerequisites are met:

- The switch is onboarded to the Mist cloud using the claim code or activation code.
- The switch is running a CloudX-supported Junos version. For more information, see ["Juniper CloudX Overview" on page 21](#).
- The DHCP server is able to hand out the proxy server information (via Option 43) and other elements such as IP Address, DNS, and default route.
- The switch can reach the HTTP proxy server over an IP network.
- The HTTP proxy server can redirect traffic to the Mist cloud. This [example](#) shows how to configure the proxy server:

To connect a switch to Mist cloud via a dynamic proxy server:

1. Power on the switch.
2. Connect the switch to the uplink (via OOB or in-band port).  
The switch sends a DHCP Discover message and accepts the Offer message along with DHCP proxy server information sent via Option 43. The switch stores the proxy server information at `/var/etc/phc_vendor_specific_info.xml`. The switch reaches out to the proxy server during the ZTP boot-up process and connects to the Mist cloud via HTTP proxy server.
3. Log in to the switch and verify the connectivity to the Mist cloud by using the following CLI command:  

```
show system connections | grep port used for connectivity between switch and proxy offered by DHCP
```
4. In case the switch does not connect to the cloud, collect logs from the following files on the switch and create a support ticket:  
`/var/log/mcd.log`, `/var/log/messages` and RSI  
For more information, see [How to collect logs and files from standalone and Virtual Chassis/VCF devices](#).

## Connect a Switch to Mist Cloud via a Static Proxy Server

If a switch cannot receive the proxy information via DHCP, you can configure it with a static proxy server through which the switch can connect to the Mist cloud. In this case, the DHCP server does not hand out the proxy server information via Option 43.

Before you connect the switch to the Mist cloud via a static proxy server, ensure that the following prerequisites are met:

- The switch is onboarded to the Mist cloud using the claim code or activation code.
- The switch has the configuration management option enabled in Mist. If not, you will need to use the switch CLI to configure the proxy server.
- The switch is running the Junos version 21.4R3-S4, 22.4R2-S1, or above.
- The local DHCP server is able to hand out IP address, DNS, default route, or statically defined route on the switch. This process involves staging the switch before establishing the cloud connectivity. If this prerequisite is met, the switch will be able to reach the HTTP proxy server over an IP network.

To connect a switch to Mist cloud via a static proxy server:

1. Log in to the Mist portal (manage.mist.com).
2. Click **Organization** > **Site Configuration** > *site-name* to navigate to the site where the switch is onboarded.
3. On the Site Proxy tile of the site configuration page, configure the proxy information, as shown below:

The screenshot shows the Mist portal configuration page for a site. The 'Site Proxy' section is highlighted with a red box. It contains the following fields and options:

- Proxy URL:** A text input field containing the value `http://john@test.com:1234`.
- Override Organization Setting:** A checkbox that is checked.

Other visible sections include:

- Switch Management:** Root Password field, Switch Proxy (Enabled/Disabled).
- WAN Edge Management:** Root Password field.
- Mist Edge Management:** Override Organization Settings checkbox, FIPS (Enabled/Disabled).
- Session Smart Conductor:** Addresses field.
- Auto-Prevent Clients:** Checkbox.
- AP Config Persistence:** Enable checkbox.
- AP Uplink Monitoring:** Enable checkbox.
- WootCloud Single Sign-on:** Enable/Disable radio buttons.
- Juniper ATP:** Enable/Disable radio buttons.
- WAN Speed Test Scheduler:** Override Organization Setting checkbox.
- Marvis Minis:** Override Organization Setting checkbox, Disable Marvis Minis checkbox.
- Custom URLs:** Add Custom URLs button.
- VLANs:** Table with columns for VLANs and URL.
- Excluded VLANs:** Add button.

4. Before you connect the switch to the Mist cloud for the first time, edit the `/var/etc/phc_vendor_specific_info.xml` file as shown below:

```
cat /var/etc/phc_vendor_specific_info.xml
<vendor-specific-information>
    <vendor-info>
        <interface>irb.0</interface>
```

```

        <proxy-server>http://192.168.1.5:3128</proxy-server>
    </vendor-info>
</vendor-specific-information>

```

Editing the `/var/etc/phc_vendor_specific_info.xml` file is a one-time activity. You can skip this step if you can stage the switch in a non-proxy environment to connect to the cloud. If the switches are staged, they can gather the proxy information from the Mist cloud (which you configured in the previous step).

When you complete the above steps, the switch will be able to reach the proxy server during the ZTP process.

5. If the switch is not connecting to the proxy server during the ZTP process, flap the uplink port to force the switch to connect to the proxy server.
6. Log in to the switch and verify the connectivity to Mist cloud by using the following CLI command:  
`show system connections | grep port used for connectivity between switch and proxy`
7. In case the switch is not connecting to the cloud, collect logs from the following files on the switch and open a support case with Juniper support:  
`/var/log/mcd.log`, `/var/log/messages` and RSI  
 For more information, see [How to collect logs and files from standalone and Virtual Chassis/VCF devices](#).

## Configure the System Log

---

### SUMMARY

Send system log messages to files, remote destinations, user terminals, or to the system console.

---

Junos OS generates system log messages (also called *syslog messages*) to record events that occur on a switch, including the following:


- Routine operations, such as creation of an Open Shortest Path First (*OSPF*) protocol adjacency or a user login to the configuration database
- Failure and error conditions, such as failure to access a configuration file or unexpected closure of a connection to a peer process

- Emergency or critical conditions, such as router power-down due to excessive temperature

For the switches that you manage in the Juniper Mist portal, you can configure syslog in the switch settings.

To configure the system log:

1. Go to the appropriate configuration page:
  - To configure an organization-level switch template—From the left menu, select **Organization > Wired > Switch Templates**.
  - To configure site-level switch settings—From the left menu, select **Site > Wired > Switch Configuration**. Click the site that you want to configure.
2. Under **All Switches Configuration**, find the **Syslog** section.
3. Click **Enabled**.



4. Click a tab, and then enter the settings.

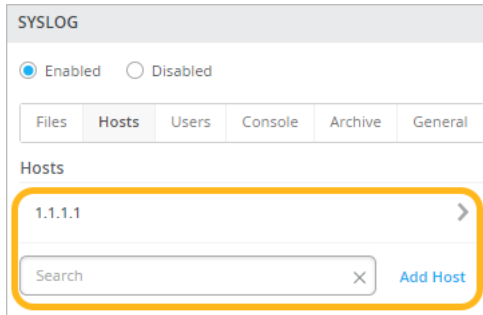
Tabs:

- **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 a time format, routing instance, and source address for the log messages.

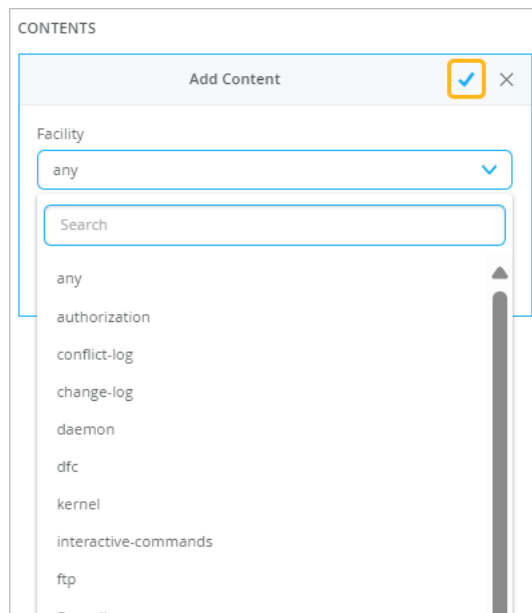
Tips for **Files**, **Hosts**, and **Users**:

- Click the **Add** link, or click an existing file, host, or user to edit.

For example, on the Hosts tab, click a host to edit, or click **Add Host**.



- In the pop-up window, add/enter the information. Or, if you're editing an existing record, you can click the delete button to remove the selected file, host, or user.
- Use the Contents section of the pop-up window to specify the messages to be sent. Click **Add Content**, and then select the **Facility** and **Level**. Click the checkmark to save these options. Repeat as needed to identify all messages to be sent.



- In the Archive section (for Files only), set the threshold for archiving by specifying the number of files and/or the total file size.
  - To save the settings in the pop-up window, click **Add** or **Save** at the bottom of the window.
5. When you're finished with all your changes, click **Save** at the top-right corner of the configuration page.

# 3

CHAPTER

## Switch Dashboards

---

[Switch Metrics](#) | 148

[Switch Details](#) | 149

[Switch Utilities](#) | 161

[Wired Clients](#) | 163


---

# Switch Metrics

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

To view switch metrics, click **Switches** on the left navigation pane on the Mist portal. Metrics for which there is no data appear grey (usually, because the feature has not been configured).

**Figure 5: Switch Metrics on the Switches Page**



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

- **Switch-AP Affinity**—This indicator shows the weighted percentage of the switches for which the number of APs connected exceeds the threshold configured. By default, the Switch-AP Affinity threshold is set to 12 APs per switch. You can configure a threshold value for the number of APs per switch to be considered in the Switch-AP Affinity metric calculation. To configure the AP threshold, click the Switch-AP Affinity indicator, and then click the hamburger icon on the right of the Switch-AP Affinity section.
- **PoE Compliance**—This indicator shows the percentage of APs that have the required 802.3at power. PoE compliance is impacted when the APs draw more power from the switch than what is allocated.
- **VLANs**—This indicator shows the percentage of APs for which all the wired VLANs are active. Click this indicator to view the list of switches and APs that have inactive or missing VLANs, along with the port information.
- **Version Compliance**— This indicator shows the percentage of switches that have the same Junos software version (per switch model). To achieve 100 percent version compliance, you must ensure that all the switches in your site run the same Junos version per switch model.
- **Switch Uptime**—This indicator shows the percentage of time a switch was up during the past seven days, averaged across all switches. Click this indicator to view the list of switches that had less than 100 percent uptime during the past 7 days.

You can click each of the switch metric indicators to get a filtered view and quickly access each device dashboard.

For more details on the switch metrics, watch the following video:



Video: [Switch Metrics Overview](#)

## Switch Details

### IN THIS SECTION

- [Finding the Switch Details Page | 149](#)
- [Front Panel | 150](#)
- [Live Port Traffic and Device Processes | 154](#)
- [Port List | 155](#)
- [Switch Insights | 156](#)
- [Metrics | 160](#)
- [Properties | 160](#)
- [Statistics | 161](#)
- [Switch Utilities | 161](#)
- [Switch Configuration | 161](#)

The switch details page is your ultimate go-to place on Mist for everything you need to know about a switch. You can view the status of each port and the statistics of the devices connected to the switch, access the switch configuration, review and modify the configuration, and track how the switch is performing against key metrics that matter to you.

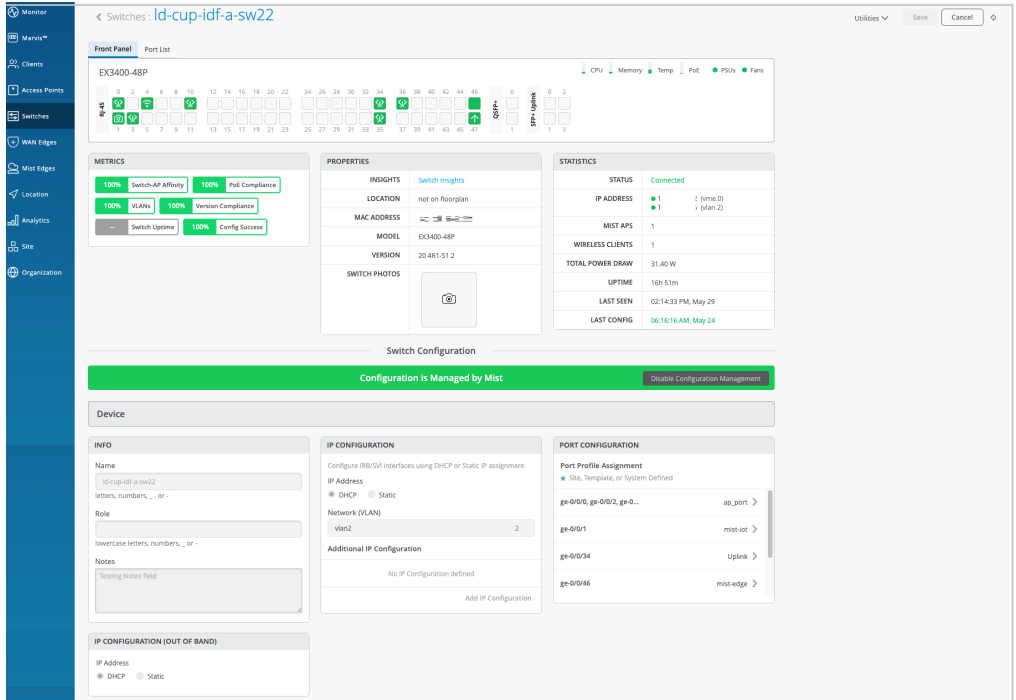
The switch details page also has the tools that help you with switch testing and troubleshooting.

### Finding the Switch Details Page

From the left menu, select **Switches**. Then click a switch in the list.

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






## 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 12: Port Icons and Their Descriptions












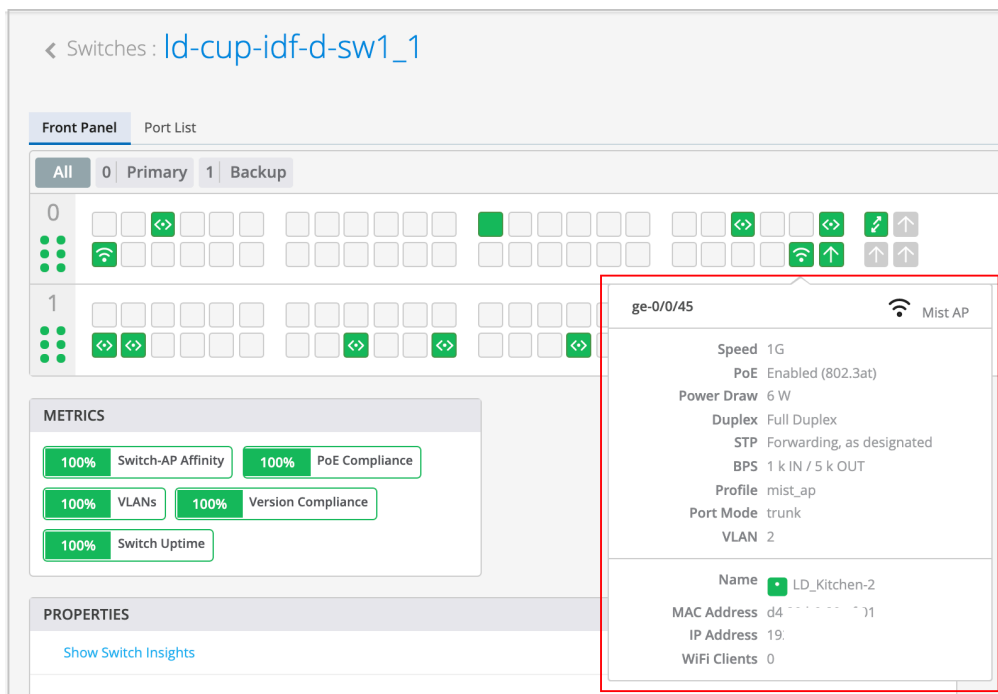
Port Icon	Description
	The port is empty. No device is connected.
	A wired client is connected.
	A wired client is connected (trunk port).
	A Mist AP is connected.
	A camera is connected. This icon applies only to Verkada cameras.
	Virtual Chassis port (VCP). A member device is connected.
	<p>The port is up.</p> <p>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.</p>
	The port is empty, with active alerts.
	A wired client is connected, but the port has active alerts.

Table 12: 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:



The screenshot shows a network management interface for a switch named 'ld-cup-idf-d-sw1\_1'. The 'Front Panel' view displays a grid of ports. A tooltip is shown for port 'ge-0/0/45', which is connected to a 'Mist AP'. The tooltip details the following information:

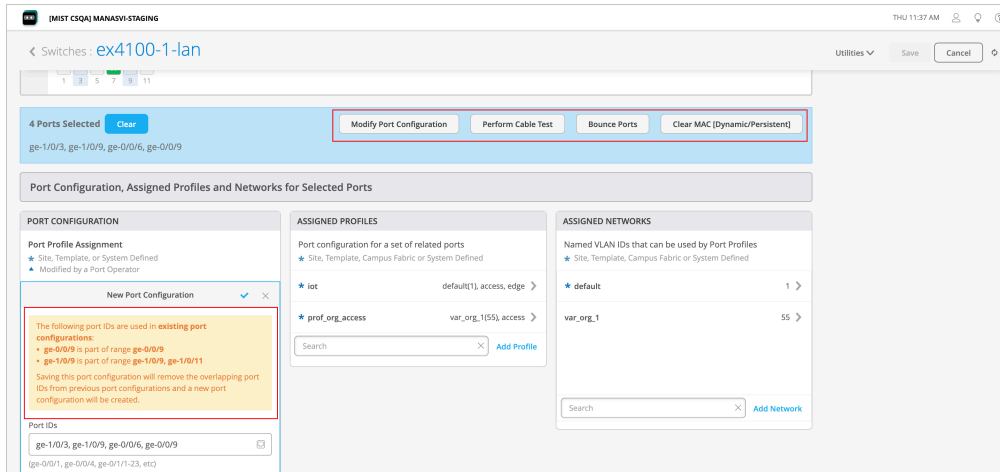
- Speed: 1G
- PoE: Enabled (802.3at)
- Power Draw: 6 W
- Duplex: Full Duplex
- STP: Forwarding, as designated
- BPS: 1 k IN / 5 k OUT
- Profile: mist\_ap
- Port Mode: trunk
- VLAN: 2
- Name: LD\_Kitchen-2
- MAC Address: d4 ... .. 1
- IP Address: 19
- WiFi Clients: 0

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.

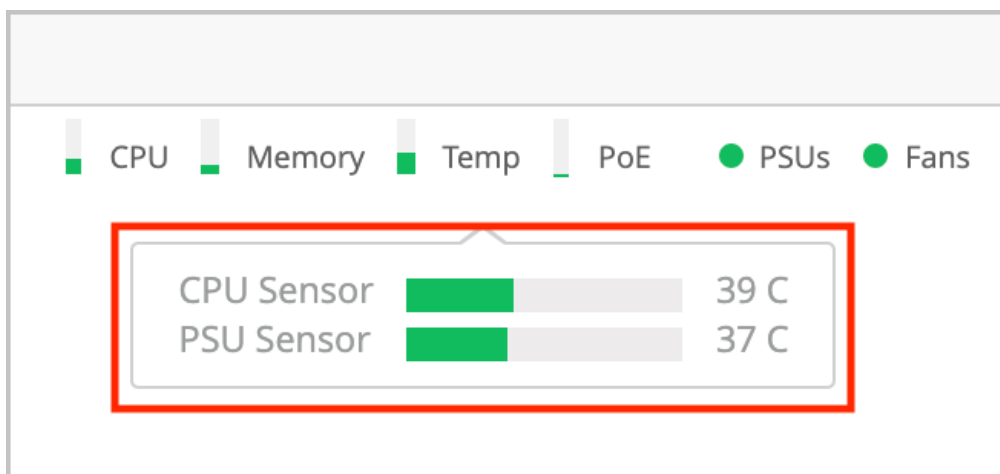
- If you select only a single port, the configuration page additionally displays port Statistics and Wired Client insights. From the Wired Client insights tile, you can access the connected devices.

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.



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.

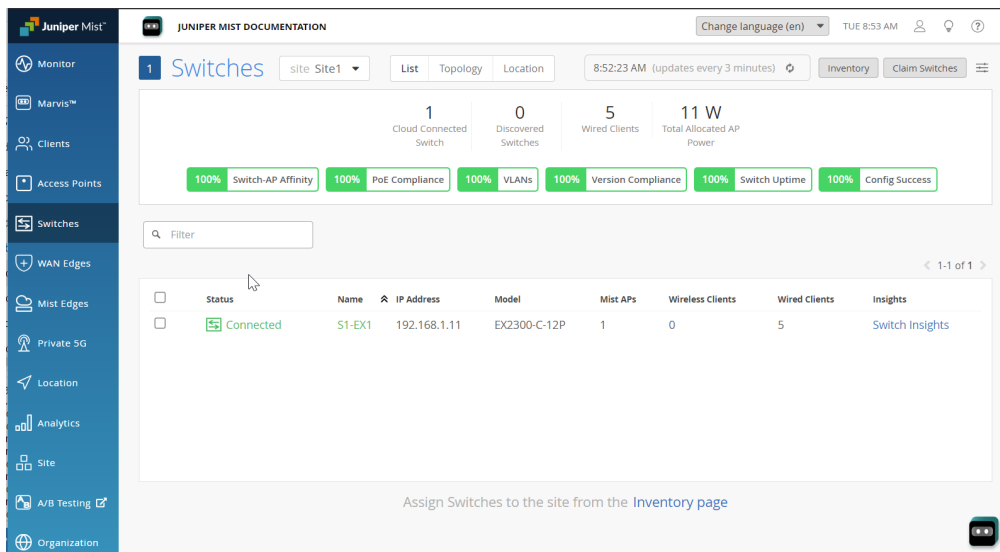


All in all, this detailed view and the additional features make managing your switch a breeze.

## Live Port Traffic and Device Processes

Users with Super User or Network Admin privileges can get real-time statistics for a given switch interface by selecting the switch and then drilling down on the port. Traffic statistics include input and output, L2 and L3 errors, and BUM traffic, and are available on all EX series and QFX series switches with an active Wired Assurance license.


Figure 6: Live Traffic Counters



### Viewing Switch Processes

You can get real-time statistics for the processes running on a given switch by selecting the switch and then drilling down to the switch Insights page, then clicking the **View Live Process Detail** button.

Figure 7: Live Process Detail

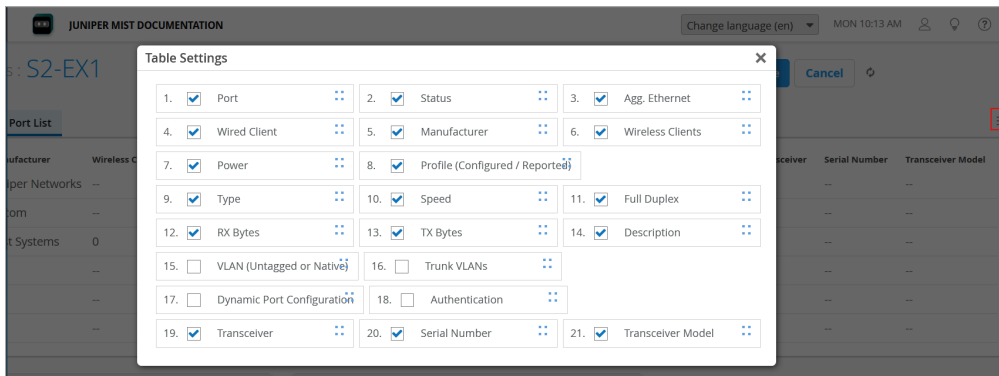


## Port List

If you want to see a list of all the ports, just click the **Port List** tab right beside the **Front Panel** tab.

In this Port List view, you'll find a wealth of information about each port. You can see the port name, its status, the clients connected to it, the power draw, port profile, port mode, port speed, and even the amount of data transmitted and received. You can also access each connected client by clicking the client name hyperlink.

By default, Port List displays a select set of columns. You can add more columns to the view from the list of available columns. To do that, go to Table Settings by clicking the hamburger menu in the upper right of the page. From the Table Settings page, select a check box to display the corresponding column (see the image below).



## Switch Insights

If you want to gain valuable insights into your switch, the events such as switch configuration changes, performance, and connected endpoints, visit the Switch Insights page by clicking the **Switch Insights** hyperlink on the **Properties** tile.



**NOTE:** You can also go to the Switch Insights page by using the hyperlink on the main Switches page.


Switch Insights gives you a bird's-eye view of all the switch events that have taken place. It's like a log of configuration changes, software updates, and system alarms.

Switch Insights also lets you track some important metrics like CPU and memory utilization of the switch. You can also dive into details about BGP neighbors (applicable to campus fabrics). You can also view traffic patterns, port errors, and power draw at the switch or port level.

The key sections on the Switch Insights page are the following:

- **Switch Events**—Provides a log of all the switch events reported. You can filter good, neutral, and bad events.
- **Table Capacity**—Provides the following indicators:
  - **MAC Address Table**—Displays the percentage of the MAC address table capacity used. The MAC address table contains MAC Address-interface bindings associated with each VLAN.
  - **ARP Table**—Displays the percentage of Address Resolution Protocol (ARP) table capacity used. The ARP table contains the learned MAC Address-IP bindings of the devices connected to the network.
  - **Route Summary**—Displays the percentage of routing table capacity used.

- EVPN Database—Available for switches that are part of an EVPN topology.




You can click the **Search Entries** button under each metric to open a shell view in a new window where you can search for entries after specifying filters. You also have the option to refresh and clear the entries displayed. Clicking **Refresh** on the upper right of the window provides a continuous display of the entry every three seconds for a total of 30 seconds. To stop the refresh before the 30-second timer is complete, close the window or click another table. Clicking the **Clear Entry** button, which is available only for MAC and ARP table, clears the respective entry from the table. You also have the option to clear the buffer on the screen by clicking **Clear Screen** at the lower left of the window.

- Switch Charts—A set of charts providing insights into key metrics.

Use the drop-down menu at the top of this section to select all ports or a specific port.

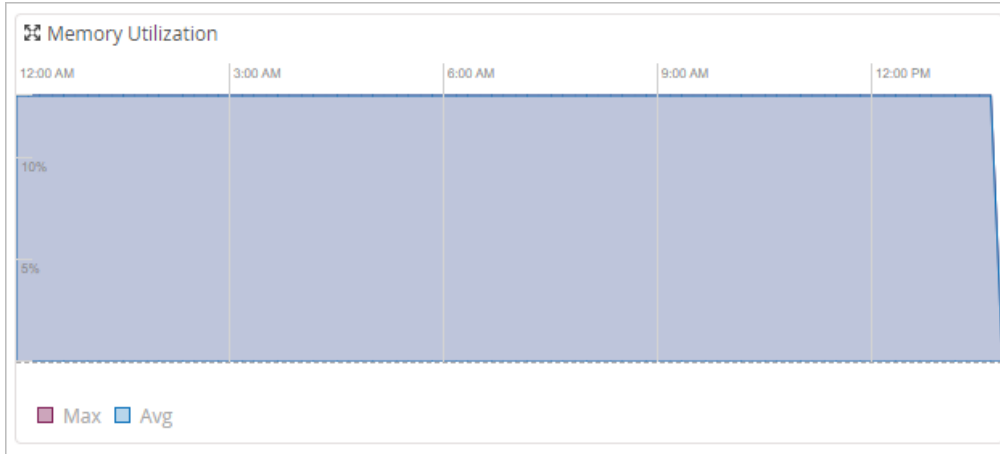
Charts include:

- CPU Utilization—Percent of CPU in use across the selected time period. Hover your mouse over any point in time to see details below the chart.




- Memory Utilization—Percent of memory in use across the selected time period. Hover your mouse over any point in time to see details below the chart.







- Bytes—Number of bytes transmitted (TX) and received (RX) across the selected time period. Hover your mouse over any point in time to see details in the white bar at the center of the chart.




- Data Rate—Transmit (TX) and receive (RX) rate (in mbps) across the selected time period. Hover your mouse over any point in time to see details in the white bar at the center of the chart.




- TX/RX Packets—Number of packets transmitted (TX) and received (RX) across the selected time period. Color coding shows the different types of packets (multicast, broadcast, and unicast). Hover your mouse over any point in time to see details in a pop-up description.



- Port Errors—Number of transmit (TX) and receive (RX) errors across the selected time period. Hover your mouse over any point in time to see details below the chart.



- Power Draw—Maximum and minimum Watts drawn across the selected time period. Hover your mouse over any point in time to see details below the chart.



- **Switch Ports**—A detailed port list. You can see all the ports along with information about the connected endpoints.



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 148.

## Properties

The Properties tile displays the switch properties that include the device location, MAC address, device ID, device model, and the Junos version on the device. You can access Switch Insights from the Properties tile.

## Statistics

This tile displays the switch connection status along with other details such as IP address, number of connected APs, uptime, and configuration status. The Last Config field on the Statistics tile shows the configuration change timestamp immediately after a user makes a configuration change to the switch.

## Switch Utilities

The switch details page provides troubleshooting and testing tools to help you get to the root of any issue. To access these tools, click the **Utilities** drop-down list in the upper right corner of the page.

For more information on switch utilities, see ["Switch Utilities" on page 161](#).

## 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 31](#).

# 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. You can download a log of your remote shell session by clicking the download button in the upper right portion of the Remote Shell session.
- **Send Switch Log to Mist**—When you experience an issue with a switch, use this option to securely send the switch logs to Mist. The Juniper Mist support team uses these logs to understand the issue and provide troubleshooting support.
- **Reboot Switch**—Reboot the switch directly from the switch details page.
- **Upgrade Firmware**—Upgrade switches directly from the switch details page. Before performing an upgrade, ensure that the switch has enough storage and a stable SSH connectivity to Mist cloud. If the switch doesn't have enough space, the upgrade will fail. To free up the space, run the following command as a root user:

```
user@switch01> start shell user root
root@Switch01:RE:0% pkg setop rm previous
root@Switch01:RE:0% pkg delete old
```

For more information, see ["Upgrade Junos OS Software on Your Switch" on page 128](#).

- **Create Template**—Create a switch configuration template based on the switch configuration. See also: ["Create a Switch Configuration Template" on page 31](#).

- **Snapshot Device**—Store a recovery snapshot in the OAM (Operations, Administrations and Maintenance) volume. The recovery snapshot in the OAM volume can be used to boot the switch in case of corrupt configuration or firmware update failure.
- **Download Junos Config**—Use this tool to download the switch's Junos configuration in a text file.
- **Replace Switch**—Replace a switch without disrupting the network services on your network topology. By default, this action copies the existing switch configuration to the new switch. You can also choose to discard specific configurations that you don't want to copy to the new switch. For more information, see "[Replace a Switch](#)" on page 138.



**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 30.

- **Sync Configuration**—Sometimes, the configuration push to a switch might fail for various reasons. In such cases, you can use this option to resend the configuration to the switch. The **Sync Configuration** operation will overwrite any configuration defined on the switch manually through the CLIs.

## RELATED DOCUMENTATION

| [Switch Details](#) | 149

# Wired Clients

---

## SUMMARY

See the wired clients page for a list of devices connected to the switches in your site.

---

The wired clients page allows you to see the connected port, LLDP names, MAC addresses, VLAN IDs, and other valuable information about client devices connected to the switches in your site. The information displayed can be tuned to match your needs. This at-a-glance dashboard helps you quickly see information that would be difficult to find and correlate using the switch's CLI.

Figure 8: The Wired Clients Page

The screenshot shows the Juniper Mist 'Wired Clients' page for 'Site 1'. The page includes a search filter and a table with the following data:


Client Name	MAC Address	VLAN	Wireless Clients	Switch	Port	Insights
<> 0:1c	0:1c	1	--	ex2300c-1	ge-0/0/10	Wired Client Insights
<> 0:1c	0:1c	1	--	ex2300c-2	ge-0/0/7	Wired Client Insights
Mist_AP12-01	4c:60	1	5	ex2300c-2	ge-0/0/1	Wired Client Insights
Mist_AP12-02	d:b8	1	7	ex2300c-1	ge-0/0/6	Wired Client Insights
Mist_AP33-1	0:42	1	5	ex2300c-1	ge-0/0/2	Wired Client Insights
Mist_AP43-1	c:33	1	8	ex2300c-1	ge-0/0/0	Wired Client Insights
snx300-1	6:48	--	--	ex2300c-1	ge-0/0/11	Wired Client Insights
ex2300c-2	c:87	1	--	ex2300c-1	ge-0/0/9	Wired Client Insights

Table 1 on page 164 below shows some of the available information for the wired devices connected to the switches ex2300c-1 and ex2300c-2 on Site 1 for this test organization.

Table 13: The Wired Clients Page

Column Name	Details
Client Name	The client name is either a MAC address, the assigned name of a switch, AP, WAN edge, or Mist Edge device. The icon shown along with the name is the same one shown on the front panel display of the "Switch Details" on page 149 page.
MAC Address	Displays the MAC address for the client connected to that switch port.
VLAN	Shows the assigned VLAN for the switch port.
Wireless Clients	When an AP is attached to the switch, this column shows the count of wireless clients attached to the AP.
Switch	When there are multiple switches in a site, this column lists the name of the switch to which the wired device is attached.
Port	This lists the port on the switch to which the client devices is attached.
Insights	This is a link to the Wired Clients Insights page for the specific wired client on that row of the list.

To customize the information in the list, you can add and remove columns from the list of connected



devices by clicking the Table Settings ( ) button on the upper right part of the Wired Clients page. [Figure 9 on page 165](#) appears as shown below.

**Figure 9: The Table Settings Page**

Table Settings		
1. <input checked="" type="checkbox"/> Client Name	2. <input checked="" type="checkbox"/> MAC Address	3. <input checked="" type="checkbox"/> VLAN
4. <input type="checkbox"/> VLAN (Assigned By RADIUS)	5. <input checked="" type="checkbox"/> Wireless Clients	6. <input checked="" type="checkbox"/> Switch
7. <input checked="" type="checkbox"/> Port	8. <input type="checkbox"/> Agg. Ethernet	9. <input type="checkbox"/> Port Profile
10. <input type="checkbox"/> RX Bytes	11. <input type="checkbox"/> TX Bytes	12. <input type="checkbox"/> RX Bit Rate
13. <input type="checkbox"/> TX Bit Rate	14. <input type="checkbox"/> IPv4 Address	15. <input type="checkbox"/> IPv6 Address
16. <input type="checkbox"/> Manufacturer	17. <input type="checkbox"/> Power Draw	18. <input type="checkbox"/> Username
19. <input type="checkbox"/> Speed	20. <input type="checkbox"/> Full Duplex	21. <input checked="" type="checkbox"/> Insights

Furthermore, if you want to work with the list of devices in a spreadsheet, you can download the list by clicking the download button highlighted in [Figure 8 on page 164](#).



# 4

CHAPTER

## Virtual Chassis Configuration

---

[Virtual Chassis Overview \(Juniper Mist\) | 167](#)

[Configure a Virtual Chassis Using EX2300, EX4650, or QFX5120 Switches | 171](#)

[Configure a Virtual Chassis Using EX3400, EX4100, EX4100-F, EX4300, EX4400, or EX4600 Switches | 175](#)

[Manage a Virtual Chassis Using Mist \(Add, Delete, Replace, and Modify Members\) | 180](#)


---

# Virtual Chassis Overview (Juniper Mist)

## IN THIS SECTION

- [Mixed and Non-Mixed Virtual Chassis | 169](#)
- [Design Considerations for Virtual Chassis | 169](#)

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 175. However, the Modify Virtual Chassis workflow (see "[Manage a Virtual Chassis Using Mist \(Add, Delete, Replace, and Modify Members\)](#)" on page 180) 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 14: 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 14: 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 15: 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)

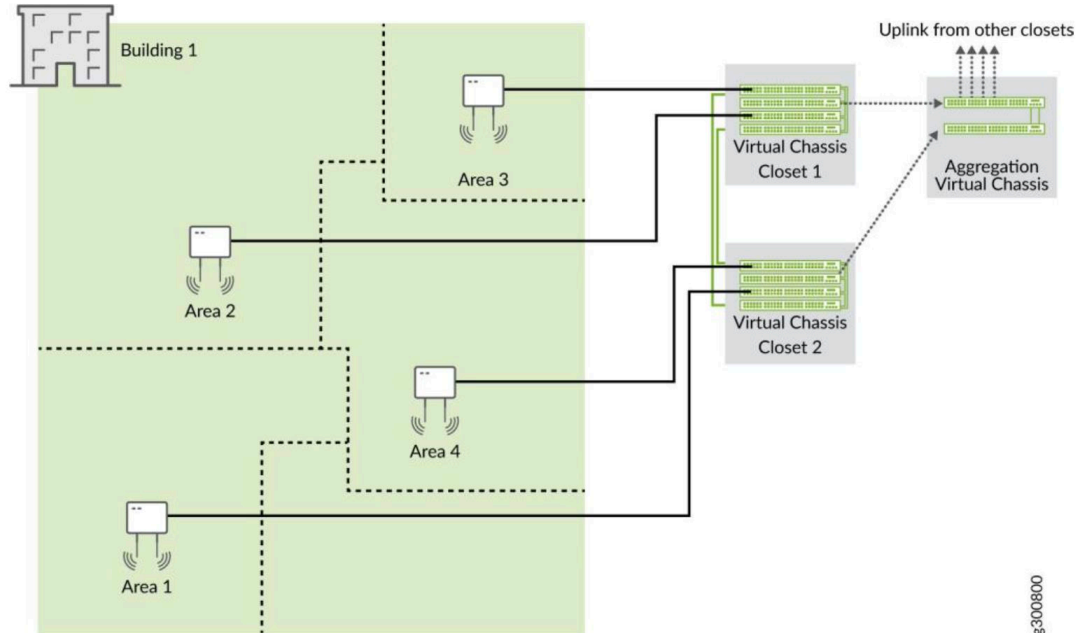
For more information about the combination of switches that a mixed or a non-mixed Virtual Chassis configuration supports, see [Understanding Mixed EX Series and QFX Series Virtual Chassis](#).

## 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 10: 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 175](#).

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 27](#). 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 128](#).

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**.

The screenshot shows the Mist Office 54 interface for the 'Switches' page. At the top, there are summary statistics: 4 Switches, 1 Mist APs, 2 Wired Clients, 0 Wireless Clients, and 18 W Total Allocated AP Power. Below these are several green status bars for 'Switch-AP Affinity', 'PoE Compliance', 'VLANs', 'Version Compliance', and 'Switch Uptime', all showing 100%. A 'More' dropdown menu is open, and the 'Form Virtual Chassis' option is highlighted with a red box. The main table below lists four switches, all with a status of 'Connected' and model 'EX2300-C-12P'.

Status	Name	IP Address	Mist APs	Wired Clients	Wireless Clients	Model
<input checked="" type="checkbox"/> Connected	Mist-ex2300-01	192.168.1.9	1	2	0	EX2300-C-12P
<input checked="" type="checkbox"/> Connected	Mist-ex2300-02	192.168.1.14	0	--	0	EX2300-C-12P
<input type="checkbox"/> Connected	Mist-ex2300-03	192.168.1.12	0	--	0	EX2300-C-12P
<input type="checkbox"/> Connected	Mist-ex2300-04	192.168.1.11	0	--	0	EX2300-C-12P

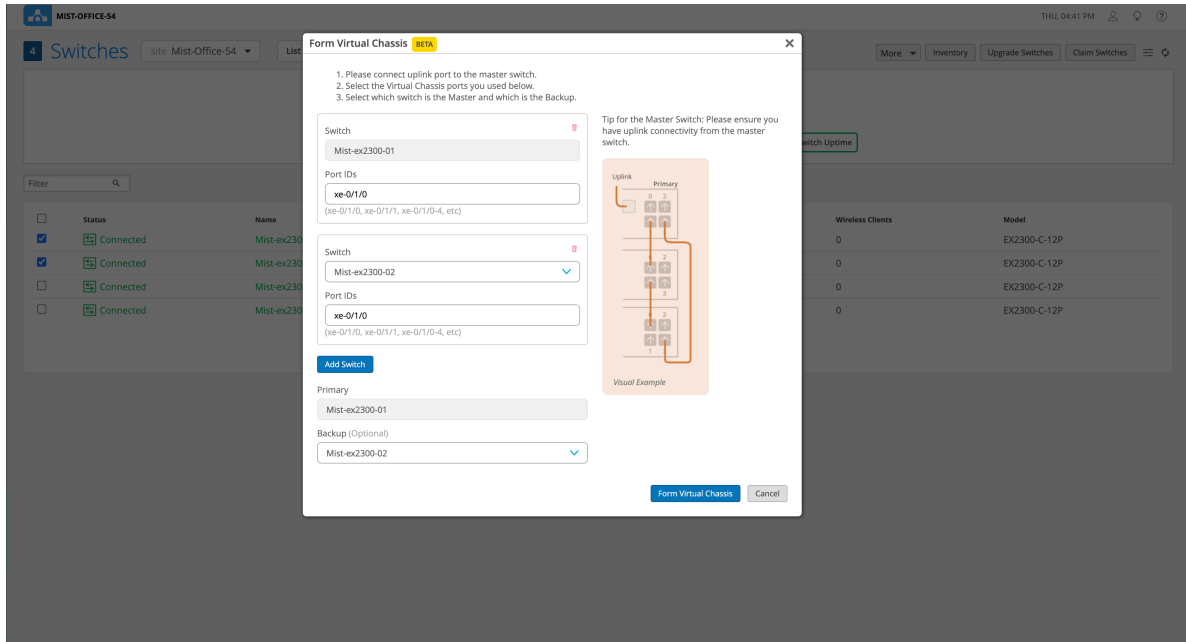


**NOTE:** You can see the **Form Virtual Chassis** option only if:

- 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.

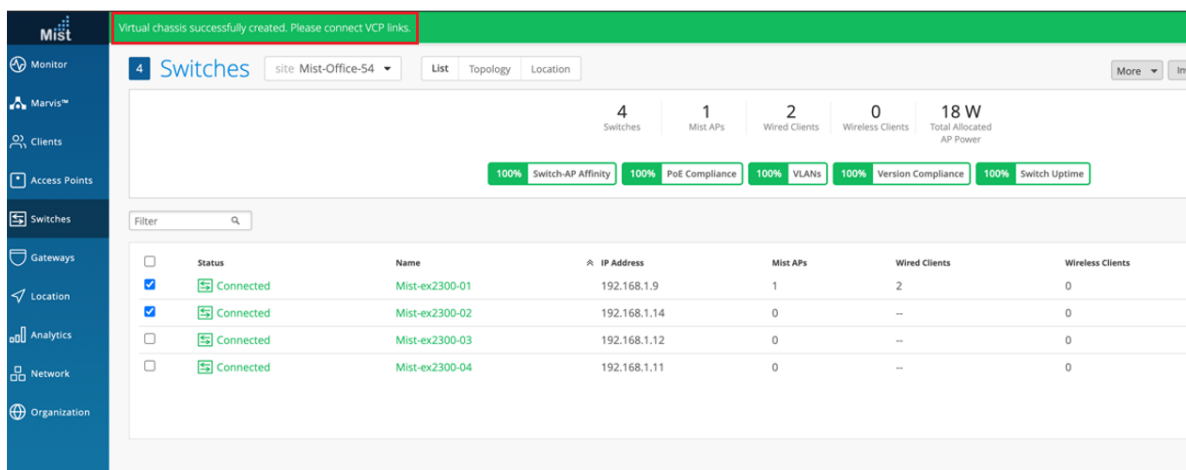
The Form Virtual Chassis window appears, as shown in the following sample picture.



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.





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**.

The screenshot shows the Mist Switches page for site 'Mist-Office-54'. The top summary bar displays: 4 Switches, 1 Mist APs, 2 Wired Clients, 0 Wireless Clients, and 18 W Total Allocated AP Power. Below this, there are five green status indicators: 100% Switch-AP Affinity, 100% PoE Compliance, 100% VLANs, 100% Version Compliance, and 100% Switch Uptime. A blue notification banner states: 'New Virtual Chassis configuration detected. It may take up to 15 minutes for the changes to show. Please refresh for the latest status.' Below the notification is a table with the following data:

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.

The screenshot shows the Mist Switches page for site 'Mist-Office-54' after successful Virtual Chassis formation. The top summary bar displays: 3 Switches, 1 Mist APs, 3 Wired Clients, 0 Wireless Clients, and 18 W Total Allocated AP Power. Below this, there are five green status indicators: 100% Switch-AP Affinity, 100% PoE Compliance, 100% VLANs, 100% Version Compliance, and 100% Switch Uptime. The table now shows only one entry for the Virtual Chassis:

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.

The screenshot displays the Mist Office 54 interface for a switch named "Mist-ex2300-01". The "Front Panel" section is highlighted with a red box, showing a "Port List" table with columns for "All", "Primary", and "Backup". Below this, there are "METRICS" for Switch-AP Affinity, PoE Compliance, VLANs, Version Compliance, and Switch Uptime, all showing 100% compliance. The "PROPERTIES" section includes a table for "Switch Insights" with columns for VC Member, Mac Address, IP Address, Model, Version, Uptime, and Status. The table shows two members: 0 (Primary) and 1 (Backup), both connected. Below the table is a "Switch Configuration" section with a green bar indicating "Configuration is Managed by Mist" and a "Disable Configuration Management" button. The bottom section contains "INFO", "PORT CONFIGURATION", and "RADIUS" 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 ["Manage a Virtual Chassis Using Mist \(Add, Delete, Replace, and Modify Members\)"](#) on page 180.

## 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 27](#).
- 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 128](#).

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 169](#).

**Table 16: 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.

**i** **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.

**i** **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.

**i** **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:

The screenshot displays the Juniper Mist interface for a Virtual Chassis named EX4400-VC. The interface is divided into several sections:

- Front Panel / Port List:** Shows three linecards (0, 1, 2) with their respective port configurations. Each linecard has a grid of ports (0-23) and a status indicator (green up arrow or green checkmark).
- METRICS:** Displays various health metrics for the Virtual Chassis, all showing 100% compliance:
  - Switch-AP Affinity: 100%
  - PoE Compliance: 100%
  - VLANs: 100%
  - Version Compliance: 100%
  - Switch Uptime: 100%
  - Config Success: 100%
- PROPERTIES:** Shows a table of Virtual Chassis members with their details.
 

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last Config
0 (Primary)	e6.0a.11.0000	ZE4322510037	10.10.10.34	EX4400-24P	22.2R3-S2.8	1h 43m	Connected	Dec 11, 2022
1 (Backup)	e6.0a.11.0001	ZE4322510023	10.10.10.34	EX4400-24P	22.2R3-S2.8	3m	Connected	Dec 11, 2022
2 (Linecard)	e6.0a.11.0002	ZE4322510045	10.10.10.34	EX4400-24P	22.2R3-S2.8	44m	Connected	Dec 11, 2022

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.

Visual Example

Available Switches ?

--
Add Switch

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.

Preprovision Virtual Chassis
Update
Cancel

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 "[Manage a Virtual Chassis Using Mist \(Add, Delete, Replace, and Modify Members\)](#)" on page 180.

# Manage a Virtual Chassis Using Mist (Add, Delete, Replace, and Modify Members)

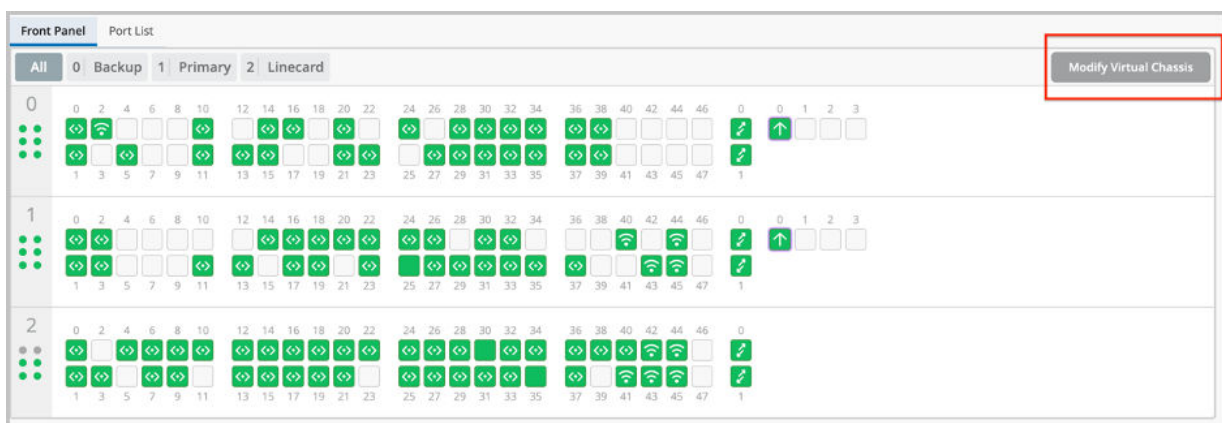
## IN THIS SECTION

- [Prerequisites | 181](#)
- [Replace a Member Switch in a Virtual Chassis | 182](#)
- [Renumber the Virtual Chassis Members | 190](#)
- [Reassign the Virtual Chassis Member Roles | 192](#)
- [Delete Virtual Chassis Members | 194](#)
- [Add a Member Switch to a Virtual Chassis | 195](#)

You can use the **Modify Virtual Chassis** option on the switch details page to manage your Virtual Chassis. The operations you can perform include renumbering and replacing the Virtual Chassis members and adding new members to a Virtual Chassis.

The Modify Virtual Chassis workflow leverages the 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 you 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

- [Replace a Non-FPC0 Member in a Virtual Chassis | 182](#)
- [Replace the FPC0 Member in a Virtual Chassis | 186](#)

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 27.

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 128 for more information.
3. Power off the new switch (the replacement switch).
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:

The screenshot shows the Mist portal interface for a Virtual Chassis. At the top, it displays the switch ID `e824a6030101`. Below this, there are tabs for "Front Panel" and "Port List". The "Front Panel" view shows three switch slots: Slot 0 (Primary), Slot 1 (Backup), and Slot 2 (Backup). Slot 0 is active, Slot 1 is offline, and Slot 2 is active. Below the front panel, there are two sections: "METRICS" and "PROPERTIES".

**METRICS**

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

**PROPERTIES**

Show Switch Insights

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last
0 (Primary)	e8:	01:01 ZE4322510037	192.168.1.101	EX4400-24P	22.2R3-S2.8	49m	Connected	De
1	f8:c	5:00 ZE4322490022	192.168.1.101	EX4400-24P	22.2R3-S2.8	--	Disconnected	De
2 (Backup)	e8:	03:01 ZE4322510045	192.168.1.101	EX4400-24P	22.2R3-S2.8	39m	Connected	De

Switch Configuration

8. Click **Modify Virtual Chassis**.

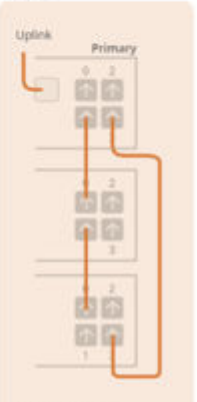
Because you have removed the VCP connection from the member switch being replaced, the Modify Virtual Chassis window displays a broken link for this member switch along with a delete (trash) icon.

**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	<b>e8:24:</b> EX4400-24P	Current Port IDs: <b>et-0/1/0, et-0/1/1</b>
1	<b>f8:c1:1</b> EX4400-24P	✖
2	<b>e8:24:</b> EX4400-24P	Current Port IDs: <b>et-2/1/2, et-2/1/3</b>

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



Visual Example

Available Switches ⓘ

e824a602e701 Add Switch

Primary

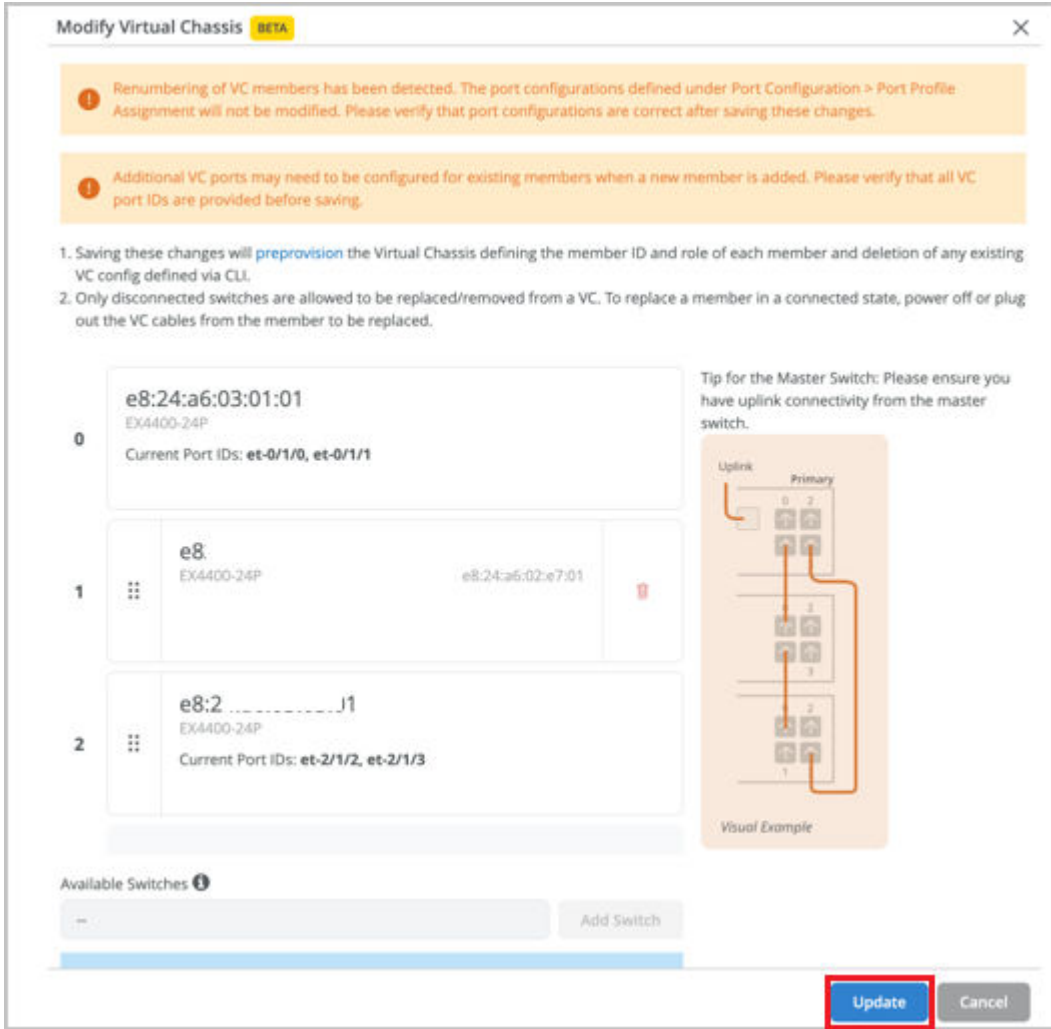
e ▼

Backup (Optional)

e8:2 ▼

Preprovision Virtual Chassis
Update
Cancel

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**.



- 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	All Event Types	All switch ports
Port Up	vcp-255/1/2	7:28:15.838 PM Dec 13, 2023						
Port Up	vcp-255/1/3	7:28:15.838 PM Dec 13, 2023						
VC Backup Elected		7:28:15.838 PM Dec 13, 2023						
VC Backup Elected		7:28:15.838 PM Dec 13, 2023						
VC Member Added		7:28:15.838 PM Dec 13, 2023						
Port Down	vcp-255/1/3	7:28:03.838 PM Dec 13, 2023						
Port Down	vcp-255/1/2	7:28:03.838 PM Dec 13, 2023						
VC Member Deleted		7:28:03.838 PM Dec 13, 2023						

Text	New adjacency to e824.a602.e701 on vcp-255/1/2.32768
Model	EX4400-24P
Version	22.2R3-S2.8

The switch details page displays the updated Virtual Chassis information.

The screenshot shows the Mist Cloud interface for a Virtual Chassis (VC) named EX4400-VC. The interface is divided into several sections:

- Front Panel:** Displays three linecards (0, 1, 2) with their respective port configurations. Each linecard has a grid of ports (0-23) and a status indicator (green dot).
- METRICS:** Shows various performance metrics, all at 100% compliance:
  - Switch-AP Affinity: 100%
  - PoE Compliance: 100%
  - VLANs: 100%
  - Version Compliance: 100%
  - Switch Uptime: 100%
  - Config Success: 100%
- PROPERTIES:** A table listing the VC members with their respective details:
 

VC Member	Mac Address	Serial Number	IP Address	Model	Version	Uptime	Status	Last Config
0 (Primary)	e8:9c:2c:00:00:00	ZE4322510037	10.10.10.34	EX4400-24P	22.2R3-S2.8	1h 43m	Connected	Dec 11, 202
1 (Backup)	e8:9c:2c:00:00:00	ZE4322510023	10.10.10.34	EX4400-24P	22.2R3-S2.8	3m	Connected	Dec 11, 202
2 (Linecard)	e8:9c:2c:00:00:00	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.

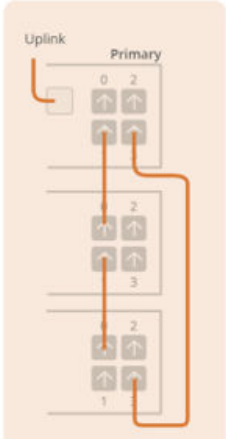
The new switch must be powered on and connected to the cloud.

### Modify Virtual Chassis BETA

- 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.
- 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		<b>e8:2</b> ..... EX4400-24P Current Port IDs: <b>et-0/1/0, et-0/1/1, et-0/1/2, et-0/1/3</b>
1	⋮	<b>f8:c</b> ..... EX4400-24P Current Port IDs: <b>et-1/1/0, et-1/1/1, et-1/1/2, et-1/1/3</b>
2	⋮	<b>e8:2</b> ..... EX4400-24P Current Port IDs: <b>et-2/1/0, et-2/1/1, et-2/1/2, et-2/1/3</b>

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



*Visual Example*

Available Switches ⓘ

e824a602e701
Add Switch

Primary

e8:2
▼

Backup (Optional)

f8:c
▼

Preprovision Virtual Chassis
Update
Cancel

4. Click **Update**.

**Modify Virtual Chassis** BETA

**!** 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.

- 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.
- Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or unplug the VC cables from the member to be replaced.

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

---

**1** **f8**  
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**

---

**3** **e1**  
EX4400-24P e8:24:a6:02:e7:01 🗑️

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

*Visual Example*

Available Switches ?

— Add Switch

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.

Update
Cancel

- 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.
- Power off the FPC0 member to be replaced. Or, remove the VCP cables from it.
- 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.
- 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.

- 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.

- Delete the member to be replaced by clicking the trash icon.

The Modify Virtual Chassis window displays a message indicating that FPC0 is required.

- 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 ✕

Remove the VC cables from the member to be replaced.

0	⋮	<b>e8:2</b> EX4400-24P Current Port IDs: <b>et-3/1/0, et-3/1/1, et-3/1/2, et-3/1/3</b>
1	⋮	<b>f8:c1</b> EX4400-24P Current Port IDs: <b>et-1/1/0, et-1/1/1, et-1/1/2, et-1/1/3</b>
2	⋮	<b>e8:2</b> EX4400-24P Current Port IDs: <b>et-2/1/0, et-2/1/1, et-2/1/2, et-2/1/3</b>

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

Uplink Primary

Visual Example

Available Switches ⓘ

-- Add Switch

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: ▼

Update Cancel

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.

- 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.

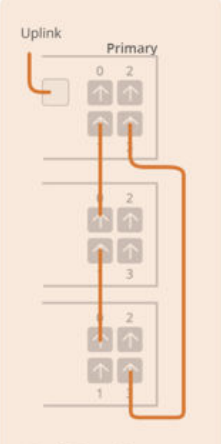
1. Saving these changes will **preprovision** the Virtual Chassis defining the member ID and role of each member and deletion of any existing VC config defined via CLI.
2. Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or plug out the VC cables from the member to be replaced.

<b>0</b>	<p>fc:50:12:10:11:11 EX3400-24P</p> <p>Current Port IDs: <b>et-0/1/0, et-0/1/1</b></p>
<b>1</b>	<p>00:c..... EX3400-48P</p> <p>Current Port IDs: <b>et-2/1/0, et-2/1/1</b></p>
<b>2</b>	<p>fc:5..... EX3400-24P</p> <p>Current Port IDs: <b>et-1/1/0, et-1/1</b></p>

Available Switches ?

--
Add Switch

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



Visual Example

Update
Cancel



#### 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. When you renumber a VC, Mist displays the following warning (as shown in the picture above): "Renumbering of VC members has been detected. The port configurations defined under Port Configuration > Port

Profile Assignment will not be modified. Please verify that port configurations are correct after saving these changes."

So, ensure that these changes are taken care of before or after renumbering the members in the Virtual Chassis.

5. After you have made the changes, click **Update**.  
The members are renumbered.

## Reassign 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 i

--
Add Switch

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.

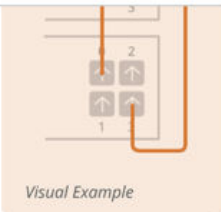
Primary

00:c5
▼

Backup (Optional)

78:5c
▼

--	
fc:	f
00	19
fc:	0
c0	30
fc:	i3
f0:	ie



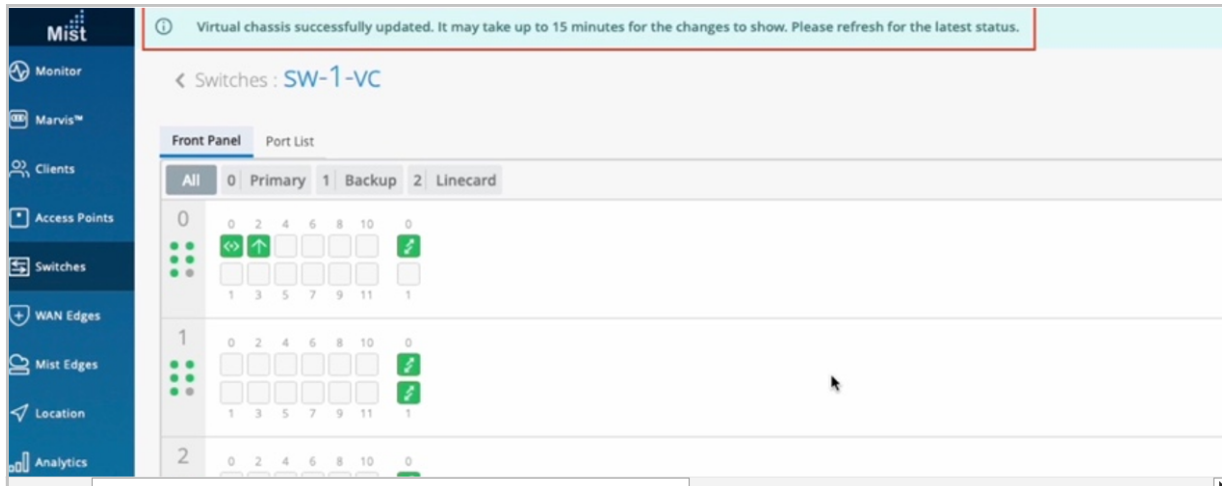
Visual Example

Update
Cancel

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:



## 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)

✕  
Delete

1
⋮

**1c:9c:8c:ba:0a:67**  
EX2300-C-12P

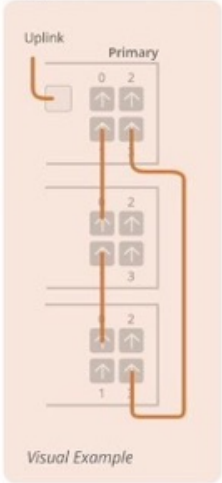
Current Port IDs: **xe-1/1/1**

VC Port IDs to Enable

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

**0c:59:9c:64:41:76**  
EX2300-C-12P

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



Visual Example

Available Switches ⓘ

EX2300-C-12P-Standalone-device-Virtual-chassis-Mist
Add Switch

Primary

f0:7c:c7:d6:6c:bd
▼

Backup (Optional)

1c:9c:8c:ba:0a:67
▼

Update
Cancel

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 27.

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 128 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.

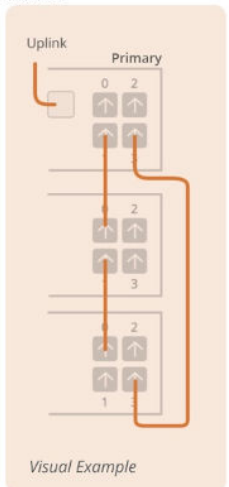
2. Only disconnected switches are allowed to be replaced/removed from a VC. To replace a member in a connected state, power off or plug out the VC cables from the member to be replaced.

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

**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)

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



*Visual Example*

Available Switches ❗

EX2300-C-12P-Standalone-device-Virtual-chassis-Mist
Add Switch

Primary

f bd
▼

Backup (Optional)

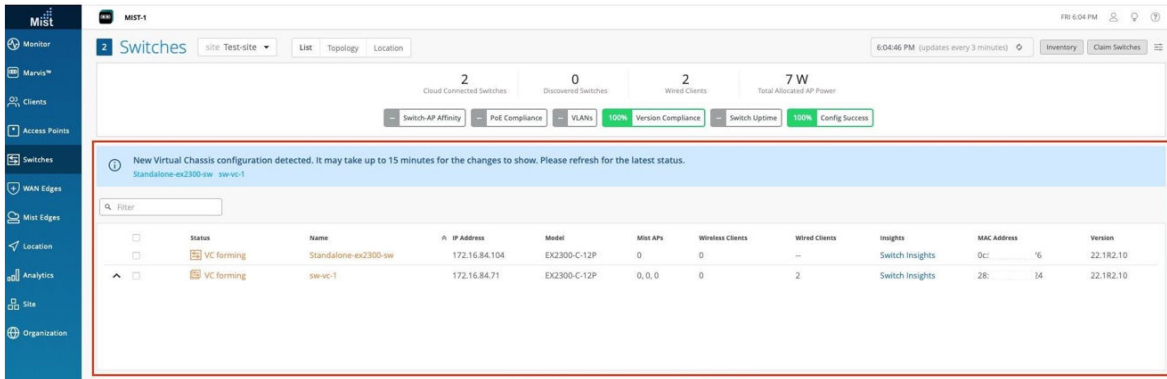
1c:9c:
▼

Update
Cancel

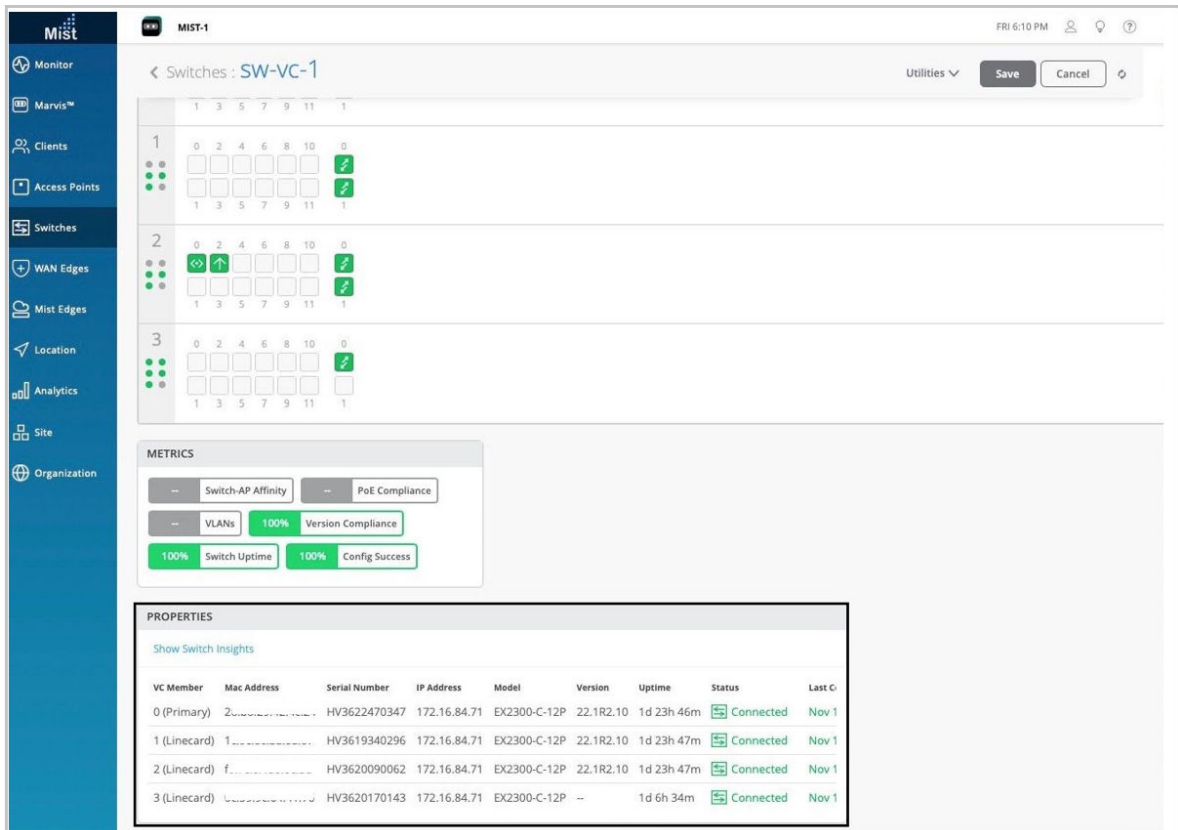
7. Specify the VC port ID for the switch, if needed (the port ID configuration applies to the EX2300, EX4650, and QFX5120 switches).
8. Click **Update**.
9. Connect the VCPs as specified on the Modify Virtual Chassis window and wait for 3 to 5 minutes for virtual chassis to be updated.



While the Virtual Chassis is forming, the switches page displays the status as 'VC Forming'.



After Mist updates the Virtual Chassis, the switch details page displays the front panel of all the three Virtual Chassis members.



# 5

CHAPTER

## Campus Fabric Configuration

---

**Which Campus Fabric Topology to Choose | 200**

How to Migrate a Traditional Enterprise Network to a Juniper Campus Fabric | 208

Configure Campus Fabric EVPN Multihoming | 213

Configure Campus Fabric Core-Distribution | 221

Configure Campus Fabric IP Clos | 231

---

# Which Campus Fabric Topology to Choose

## IN THIS SECTION

- [EVPN Multihoming for Collapsed Core | 200](#)
- [Campus Fabric Core-Distribution for Traditional 3-Stage Architecture | 203](#)
- [Campus Fabric IP Clos for Micro-Segmentation at Access Layer | 206](#)

Juniper Networks [campus fabrics](#) provide a single, standards-based Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) solution that you can deploy on any campus. You can deploy campus fabrics on a two-tier network with a collapsed core or a campus-wide system that involves multiple buildings with separate distribution and core layers.



**NOTE:** This topic lists multiple switch models that support the various campus fabric deployments. In the case of the QFX5130, all variants support the campus fabric deployments, but only one variant is supported in Mist. That variant is the QFX-5130-32CD.

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:

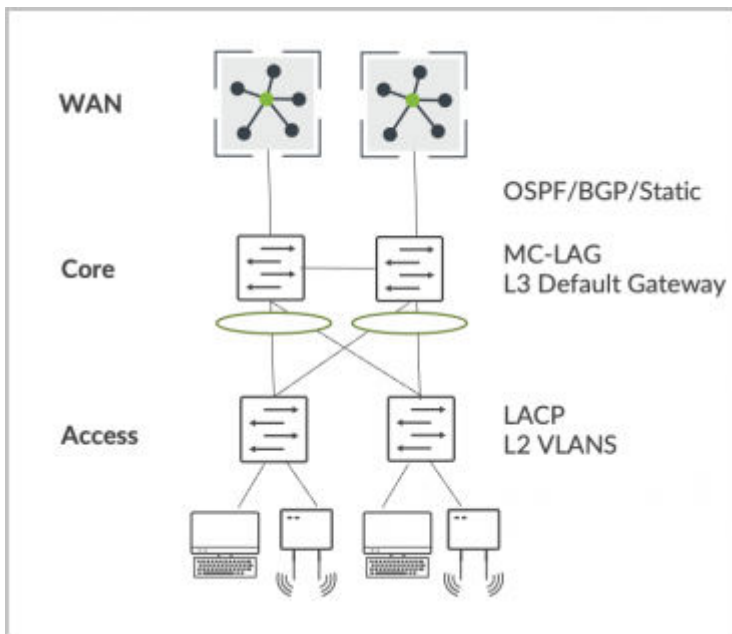
## EVPN Multihoming for Collapsed Core

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 11: Collapsed Core Topology**



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

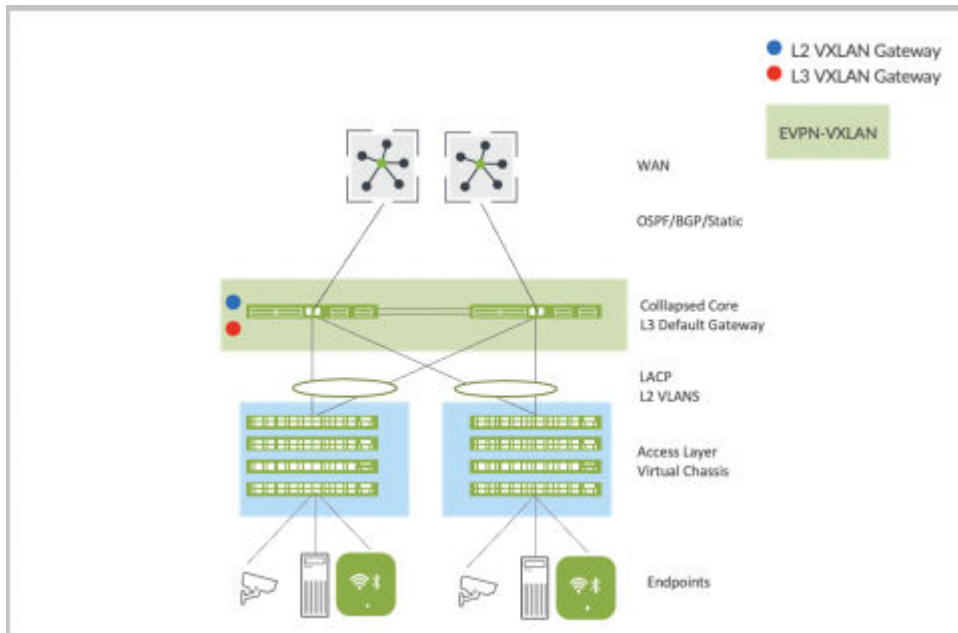
- Its proprietary MC-LAG technology requires a homogeneous vendor approach.
- It lacks horizontal scale. It supports only up to two core devices in a single topology.
- It lacks native traffic isolation capabilities in the core.
- Not all implementations support active-active load balancing to the access 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 12: 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.

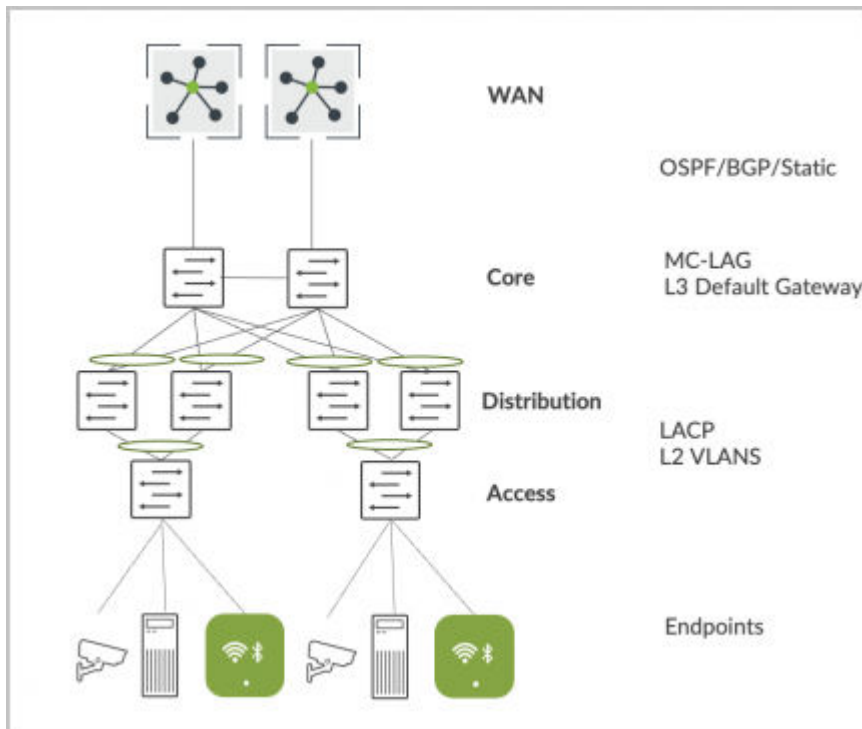
The following Juniper platforms support EVPN Multihoming:

- 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 13: 3-Stage Core-Distribution-Access Network



This traditional deployment model faces the following challenges:

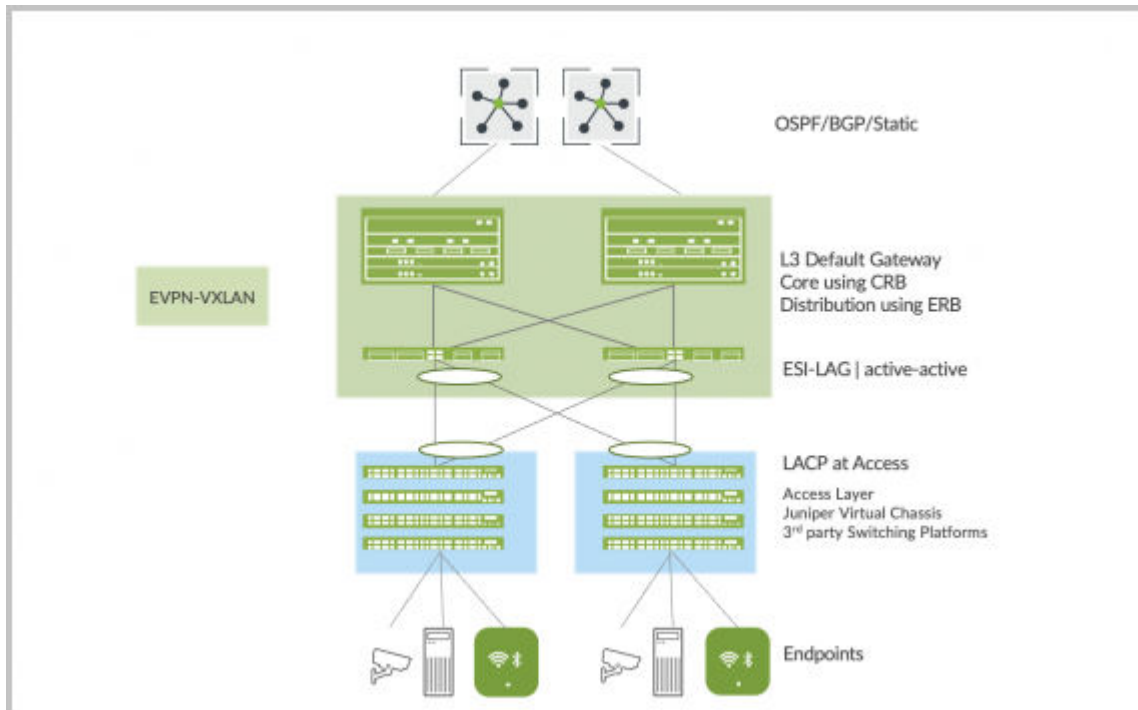
- Its proprietary core MC-LAG technology requires a homogeneous vendor approach.
- Only up to two core devices are supported in a single topology.
- Lack of native traffic isolation capabilities anywhere in this network.
- Requires STP between the distribution and access layers and potentially between the core and distribution layers. This results in sub-optimal use of links.
- Careful planning is required if you need to move the L3 boundary between core and distribution layers.
- VLAN extensibility requires deploying VLANs across all links between access switches.

The Campus Fabric Core-Distribution architecture addresses these challenges in the physical layout of a three-stage model and provides the following advantages:

- Helps in retaining your investment in the access layer. In an enterprise network, your company makes most of the Ethernet switching hardware investment in the access layer where endpoints terminate. The endpoint devices (including laptops, APs, printers, and IOT devices) plug in to the access layer. These devices use various Ethernet speeds, such as 100M, 1G, 2.5G, and 10G.
- Provides a standards-based EVPN-VXLAN framework.
- Supports horizontal scale at the core and distribution layers, supporting an IP Clos architecture.
- Provides traffic isolation capabilities native to EVPN-VXLAN.
- Provides native active-active load balancing to the access layer using ESI-LAG.
- Provides standard LACP at the access layer.
- Mitigates the need for STP between all layers.
- Supports the following topology subtypes:
  - Centrally routed bridging (CRB): Targets north-south traffic patterns with the L3 boundary or default gateway shared between all core devices.
  - Edge-routed bridging (ERB): Targets east-west traffic patterns and IP multicast with the L3 boundary or the default gateway shared between all distribution devices.

To know about more benefits of Campus Fabric Core-Distribution deployments, See [Benefits of Campus Fabric Core-Distribution](#).

Figure 14: 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

To know about more benefits of Campus Fabric IP Clos deployments, see [Benefits of Campus Fabric IP Clos](#).

The following images represents the 3-stage and 5-stage IP Clos deployment.

Figure 15: Campus Fabric IP Clos 3 Stage

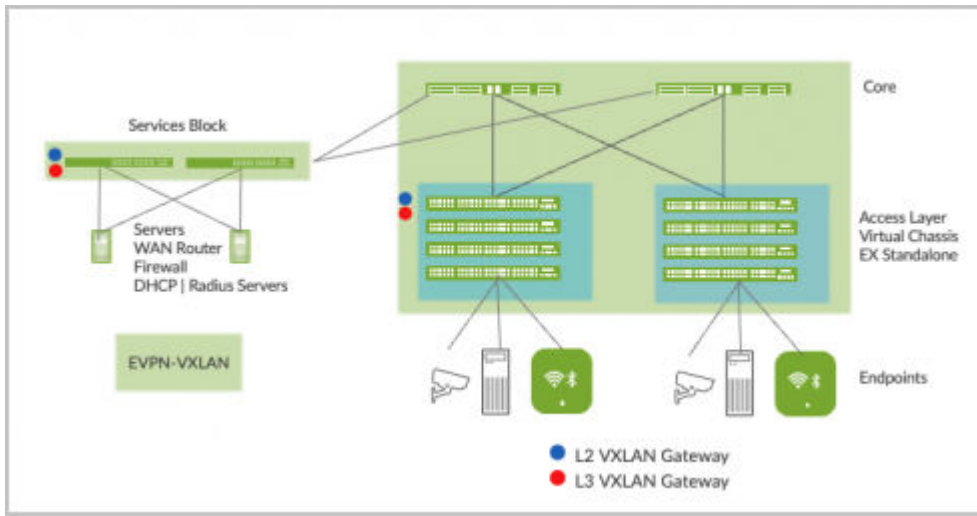
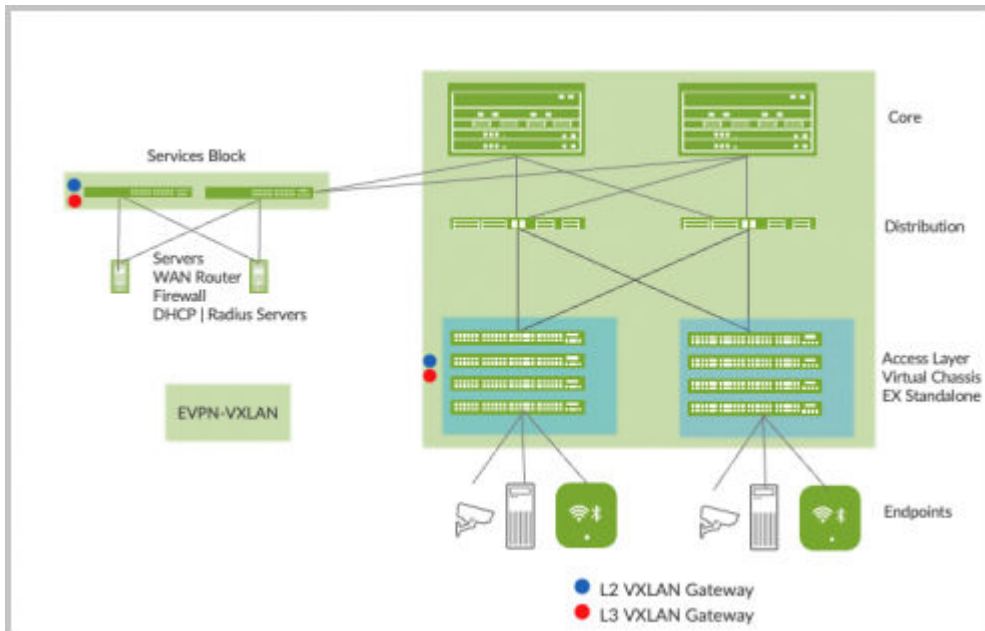


Figure 16: Campus Fabric IP Clos 5 Stage



The following Juniper Network platforms support Campus Fabric IP Clos:

- Core layer devices: EX9200, EX4400-48F, EX4400-24X, EX4650, QFX5120, QFX5110, QFX5700, and QFX5130
- Distribution layer devices: EX9200, EX4400-48F, EX4400-24X, EX4650, QFX5120, QFX5110, QFX5700, and QFX5130

- Access layer devices: EX4100, EX4300-MP, and EX4400
- Services Block devices: QFX5120, EX4650, EX4400-24X, EX4400, QFX5130, QFX5170, EX9200, and QFX10k

## How to Migrate a Traditional Enterprise Network to a Juniper Campus Fabric

### IN THIS SECTION

- Build Campus Fabric in Parallel to the Existing Network | 210
- Interconnect the Campus Fabric to the Existing Network | 210
- Migrate VLANs to the Campus Fabric | 211
- Migrate the Critical Infrastructure to the Services Block | 212
- Migrate WAN Router(s) to the Services Block | 213
- Decommission the Existing Enterprise Network | 213

This document details a strategy to migrate a traditional enterprise network-based architecture to a Juniper Campus Fabric EVPN-VXLAN architecture.

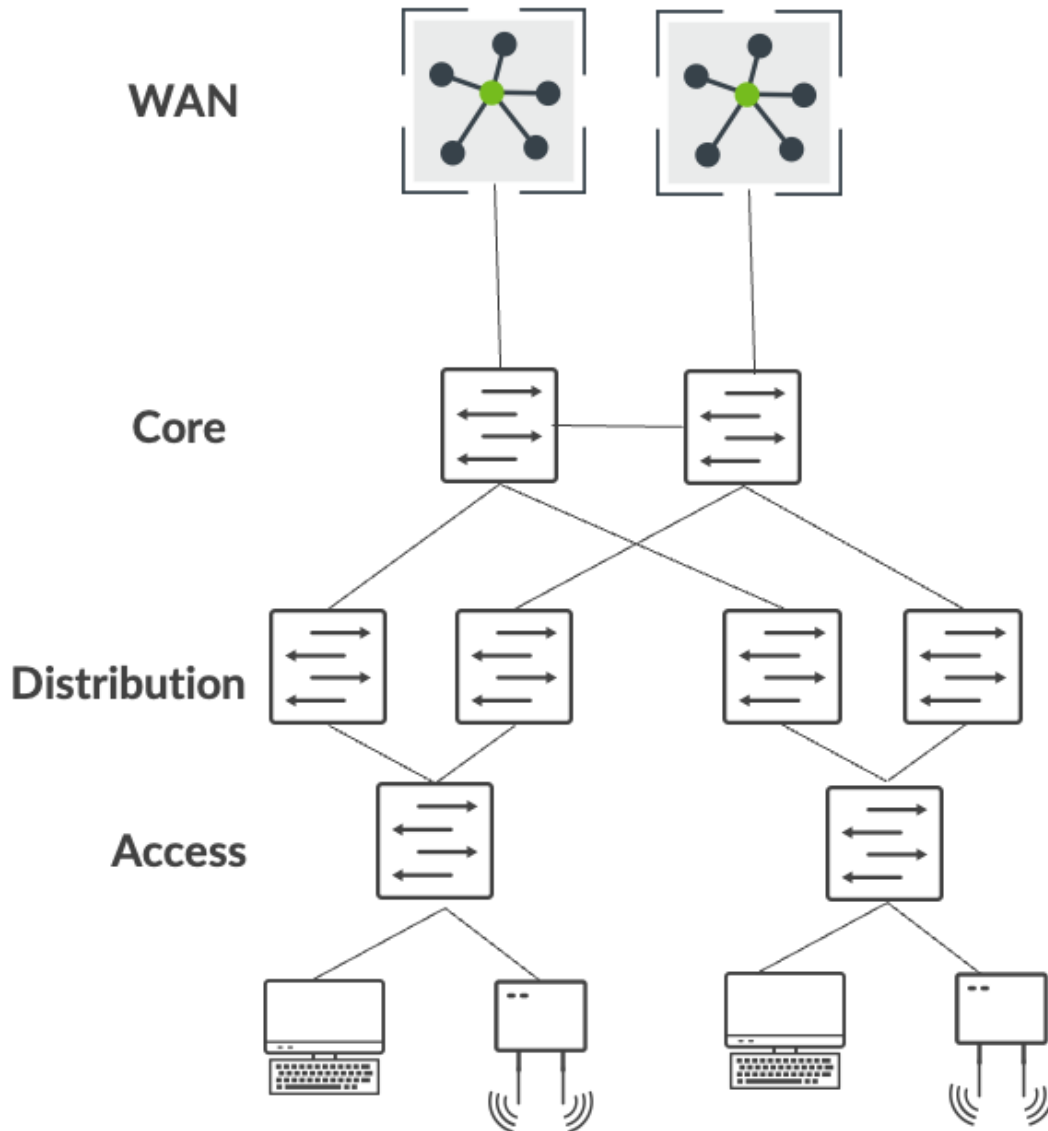
Juniper's campus fabric leverages EVPN VXLAN as the underlying technology for small, mid, and large enterprise deployments. You can build and manage campus fabric by using Mist's Wired Assurance Cloud-ready framework. For additional information on Juniper's Campus Fabric, see [Juniper Mist Wired Assurance datasheet](#).



**Video:** [Three Step Campus Fabric](#)

This migration strategy focuses on an enterprise network consisting of the traditional 3-stage architecture of access, distribution, and core. In this example, core provides layer 3 connectivity to all users, printers, access points (APs), and so on. And the core layer interconnects with dual WAN routers using standards based OSPF or BGP technologies.

Figure 17: Traditional Enterprise Network



At a high-level, migration from a traditional enterprise network to a Juniper campus fabric architecture involves the following steps:

1. Build a campus fabric architecture in parallel to the existing enterprise network.
2. Interconnect the campus fabric to the existing network using a services block.
3. Migrate VLANs one by one to the campus fabric.
4. Migrate the critical infrastructure such as DHCP server and RADIUS to the services block.

5. Migrate WAN router(s) to the services block.
6. Decommission the existing enterprise network once all the connectivity is verified.

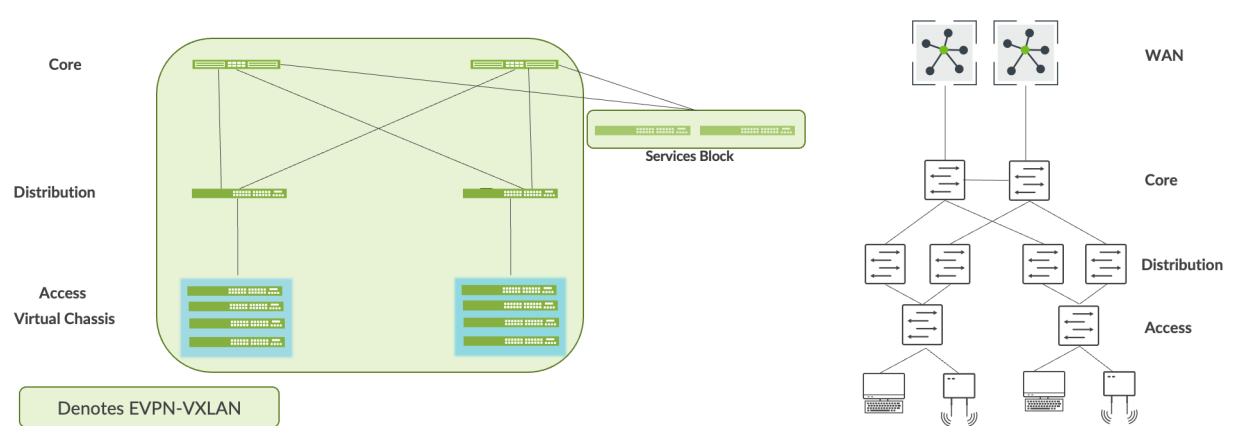
## Build Campus Fabric in Parallel to the Existing Network

As the first step, build a campus fabric by using Mist's Wired Assurance framework. This step allows you to deploy an operational campus fabric in parallel to the existing network. In this example, we choose the campus fabric IP Clos architecture because the customer has a micro-segmentation strategy deployed at the access layer. The customer has chosen the following Juniper equipment to be deployed within the Campus Fabric IP Clos architecture:

- QFX5120 switches (core layer)
- QFX5120 switches (distribution layer)
- EX4100 and EX4400 switches in Virtual Chassis mode (access layer)
- QFX5120 switches (services block)

See also: ["Configure Campus Fabric IP Clos" on page 231.](#)

**Figure 18: Co-existence of Campus Fabric with Enterprise Network**

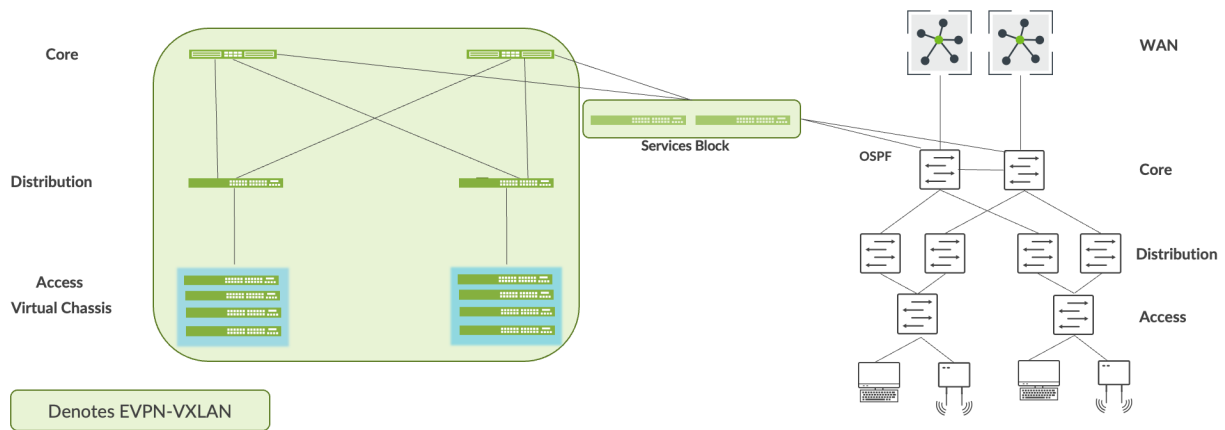


## Interconnect the Campus Fabric to the Existing Network

You can use the services block to interconnect the campus fabric with the enterprise network. You can do this by using ESI-LAG technologies at layer 2, or the standard routing protocols such as BGP or OSPF

if layer 3 connectivity is required. In this case, we interconnect the services block to the core enterprise using OSPF.

**Figure 19: Services Block Interconnects with the Core Using OSPF**



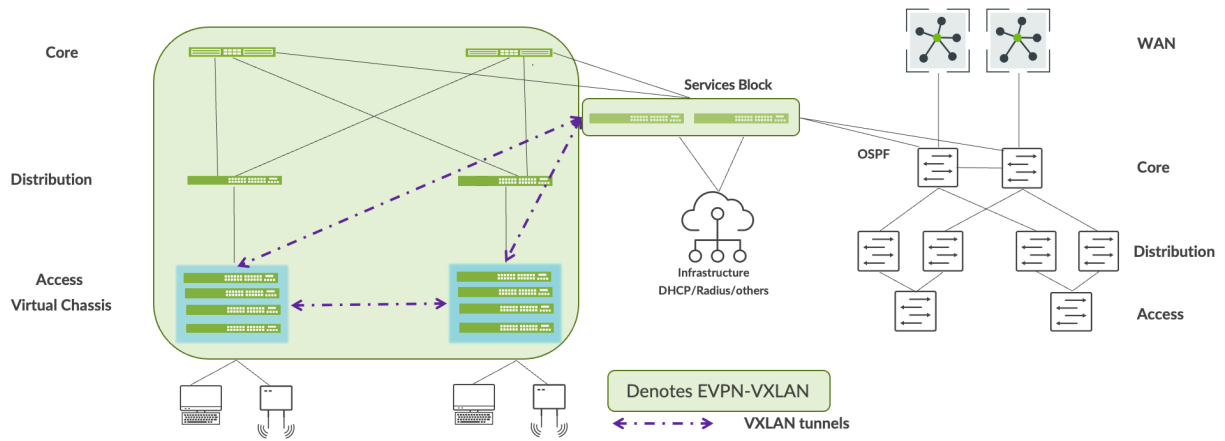
The loopback reachability between the two networks should be established through the services block. For example, the campus fabric build assigns loopback addresses to each device. By default, these addresses are part of the same subnet. OSPF should exchange these addresses with routable prefixes sent by the core layer through the services block. The end-user should verify reachability between these prefixes before moving to the next step.

## Migrate VLANs to the Campus Fabric

This process requires you to remove each VLAN and associated layer 3 interface from the enterprise network. You need to migrate all devices within the VLAN to the campus fabric and then have the end-user verify full connectivity from the devices on the migrated VLAN to the applications and devices on the enterprise network. The following summarizes this step:

- Migrate VLANs to campus fabric by disabling or removing the layer 3 subnet from the current network.
- Users and devices migrate to the access layer of the campus fabric.
- Layer 3 interconnect provides reachability on a VLAN-by-VLAN basis.
- Users and devices must validate all application reachability before moving to the next VLAN.

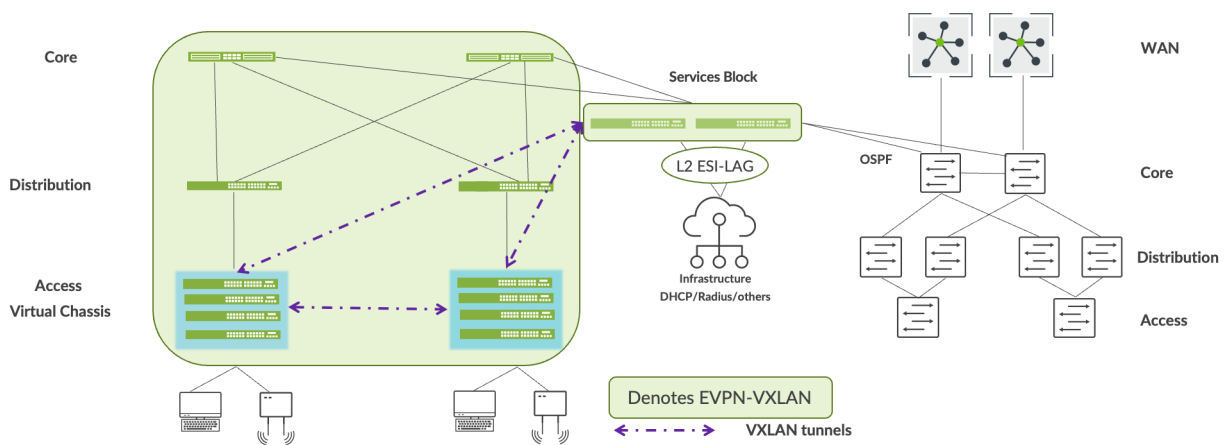
Figure 20: All VLANs and Access Devices Migrated to the Campus Fabric



## Migrate the Critical Infrastructure to the Services Block

Juniper recommends dual homing of each critical infrastructure service (such as DHCP server and RADIUS) to the services block. You can do this by using ESI-LAG technologies at layer 2, or the standard routing protocols such as BGP or OSPF if layer 3 connectivity is required. Accessibility of critical infrastructure services within the campus fabric and from the enterprise network should be verified before moving to the next step.

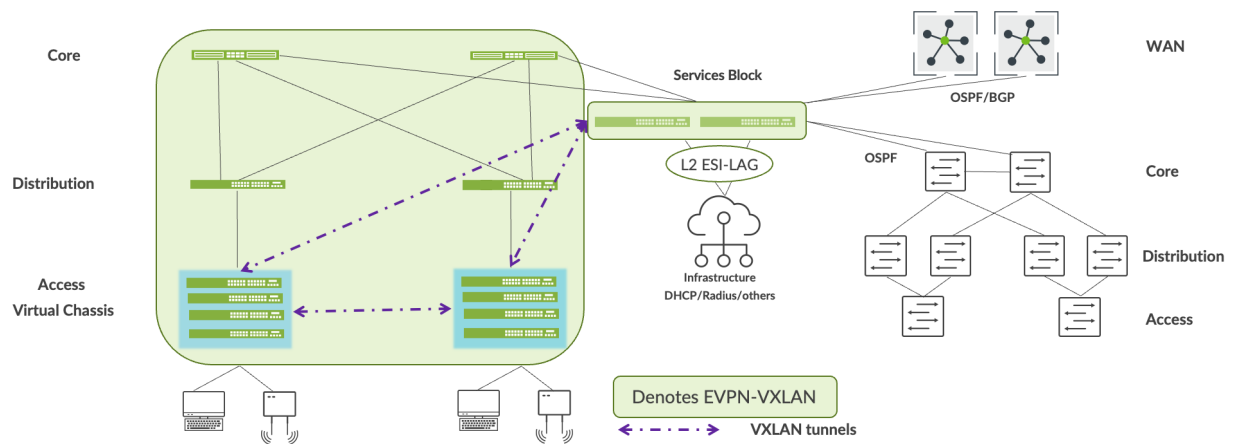
Figure 21: Critical Infrastructure Migration to the Services Block



## Migrate WAN Router(s) to the Services Block

Mist lets you connect WAN router(s) to the services block using BGP or OSPF. After WAN routers are connected to the service block, verify the accessibility of WAN services to and from the campus fabric before moving to the next step.

Figure 22: WAN Routers Migration to the Services Block



## Decommission the Existing Enterprise Network

We recommend that you keep the enterprise network up and operational for at least one week after all services and applications are reachable without issue to and from the campus fabric. After that, decommission the existing enterprise network.

## Configure Campus Fabric EVPN Multihoming

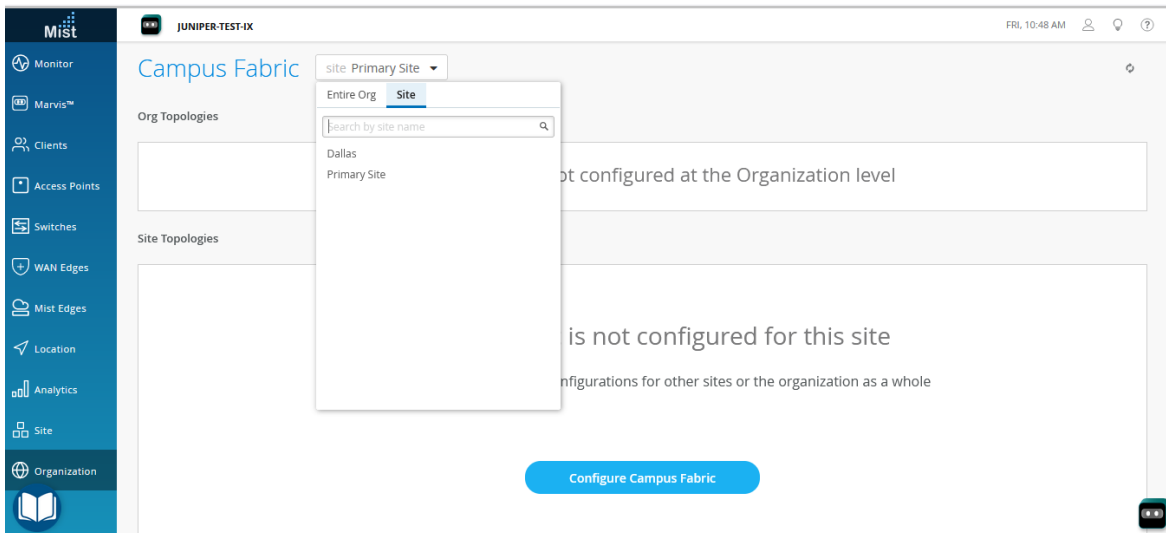
The Juniper Networks [campus fabrics](#) EVPN multihoming solution supports a collapsed core architecture. This architecture merges the core and distribution layers into a single switch. Merging these layers into a single switch turns the traditional three-tier hierarchical network into a two-tiered 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](#).



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.



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**.

**Campus Fabric Configuration** 1. Topology 2. Nodes 3. Network Settings 4. Ports 5. Confirm ← Back Continue →

## Choose Campus Fabric Topology

Choose the topology you want to construct and configure related options

**TOPOLOGY TYPE**

- EVPN Multihoming**  
Collapsed core with ESI-Lag
- Campus Fabric Core-Distribution**  
EVPN core/distribution with ESI-Lag
- Campus Fabric IP Clos**  
Campus fabric with L3 at the edge

**CONFIGURATION**

**Topology name is required**

Topology Name

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

Enabled  Disabled

**OVERLAY SETTINGS**

BGP Local AS

65000

(2-byte or 4-byte)

**UNDERLAY SETTINGS**

AS Base

65001

(2-byte or 4-byte)

Underlay

IPv4  IPv6

Subnet

2001::/64

(xxx::xxx/xx)

IPv6 Loopback Interface

fd31:db8::/32

(xxx::xxx/xx)

IPv4 Auto Router ID Subnet / Loopback Interface

172.16.254.0/23

(xxx.xxx.xxx.xxx/xx)

5. Configure the remaining settings on the **Topology** tab, as described below:



**NOTE:** We recommend that you use the default settings on this screen unless they conflict with any networks attached to the campus fabric. The point-to-point links between each layer utilize /31 addressing to conserve addresses.

- 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.
- (If you choose not to use the default settings) In the **OVERLAY SETTINGS** section, enter the following:
  - BGP Local AS**—Represents the starting point of private BGP AS numbers that Mist allocates to each device automatically. You can use any private BGP AS number range that suits your

deployment. Mist provisions the routing policy so that only the loopback IP addresses are exchanged in the underlay of the fabric.

- c. (If you choose not to use the default settings) In the **UNDERLAY SETTINGS** section, configure the following:
- **AS Base**—The AS base number. The default is 65001.
  - **Underlay**—Select an internet protocol for the underlay. Options are IPv4 and IPv6.
  - **Subnet**—The range of IP addresses that Mist uses for point-to-point links between devices. You can use a range that suits your deployment. Mist breaks this subnet into /31 subnet addressing per link. You can modify this number to suit the specific deployment scale. For example, a /24 network would provide up to 128 point-to-point /31 subnets.
  - **IPv6 Loopback Interface**—Specify an IPv6 loopback interface subnet, which is used to autoconfigure IPv6 loopback interface on each device in the fabric.
  - **Auto Router ID Subnet/Loopback Interface**—Mist uses this subnet to automatically assign a router ID to each device in the fabric (including access devices irrespective of whether they are configured with EVPN or not). Router IDs are loopback interfaces (lo0.0) used for overlay peering between devices. For new topologies, this field auto-populates a default subnet value (172.16.254.0/23), which you can modify. When you edit an existing topology, this field doesn't populate any default value. The router ID is used as an identifier when deploying routing protocols such as BGP. After you add the switch to the collapsed core layer, click the switch icon to see the associated router ID.

You can overwrite the automatically assigned router ID by manually configuring a loopback interface in the Router ID field on the Routing tile on the switch configuration page (**Switches > Switch Name**). However, if you modify the campus fabric configuration afterwards, Mist performs the automatic assignment of the router ID again, replacing the manually configured loopback interface.

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

The screenshot shows the 'Campus Fabric Configuration' interface. A modal window titled 'Select Campus Fabric Nodes' is open, displaying a table of available switches. The table has columns for Name, MAC Address, Serial, Router ID, and Model. Two switches are selected, indicated by blue checkmarks in the Name column.

Name	MAC Address	Serial	Router ID	Model
<input type="checkbox"/> f4b52ff40400	f4:b5:2f:f4:04:00	JN122EFFF5RFC		EX9204
<input type="checkbox"/> f4b52ff3f400	f4:b5:2f:f3:f4:00	JN122EFFF5RFC		EX9204
<input type="checkbox"/> DC81-DIST-SW_1_QFX5120	d8:53:9a:64:6f:c0	XH3121410895	192.168.255.11	QFX5120-48Y
<input type="checkbox"/> DC81-DIST-SW_2_QFX5120	d8:53:9a:64:b5:c0	XH3121410874	192.168.255.12	QFX5120-48Y
<input checked="" type="checkbox"/> 00cc34f3cf00	00:cc:34:f3:cf:00	ZD4422030024		EX4400-48P
<input checked="" type="checkbox"/> 00cc34f47200	00:cc:34:f4:72:00	ZD4422070133		EX4400-48P

At the bottom of the modal, there are 'Select 2' and 'Cancel' buttons. Below the modal, there are sections for 'Collapsed Core' and 'Access' layers, each with a plus sign icon and the text 'Select Switches'.

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.

9. Configure the network settings, as described below:
- a. On the **NETWORKS** tile, add networks or VLANs to the configuration. You can either create a new network or import the network from the switch template defined on the **Organization > Switch** templates page.  
To add a new VLAN, click **Create New Network** and configure the VLANs. The settings include a name, VLAN ID, and a subnet. You can specify IPv4 or IPv6 addresses for the subnet.

To import VLANs from the template:

- i. Click **Add Existing Network**.
- ii. Select a switch template from the **Template** drop-down list to view the VLANs available in that template.
- iii. Select the required VLAN from the displayed list, and click the ✓ mark.

VLANs are mapped to Virtual Network Identifiers (VNIs). You can optionally map VLANs to virtual routing and forwarding (VRF) instances to logically separate the traffic.

- b. Review the settings in the **OTHER IP CONFIGURATION** tile. This section populates the settings automatically after you specify the networks in the NETWORKS section.

- c. Optionally, configure VRF instances 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.

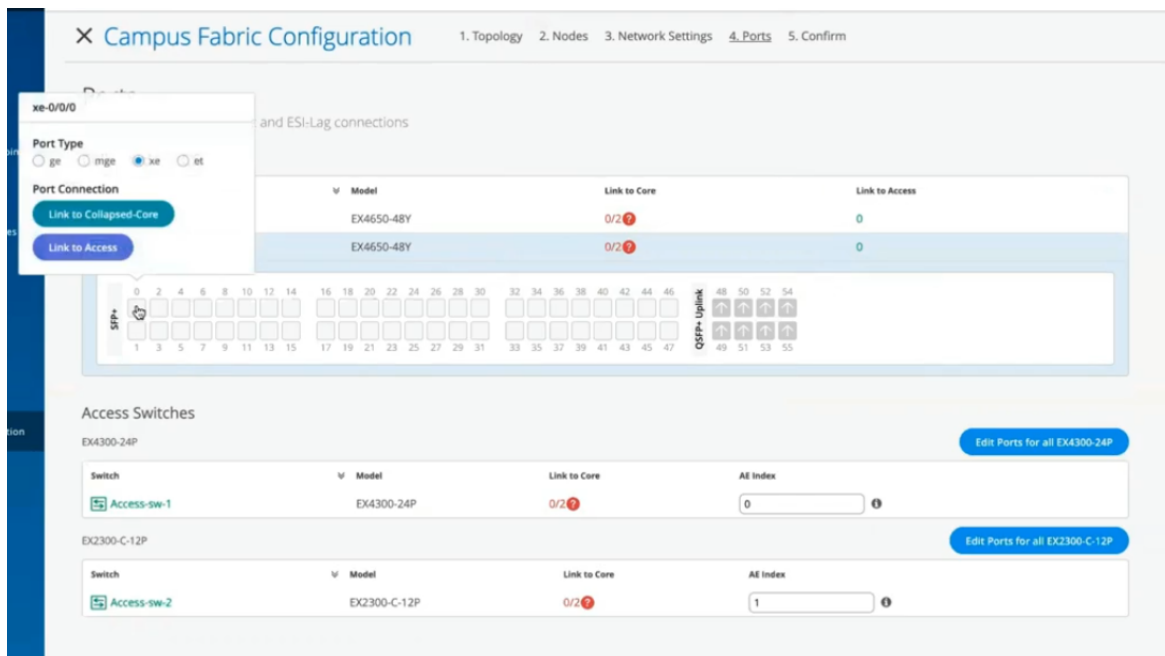
To create a VRF:

- 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. You can specify IPv4 or IPv6 addresses.
- d. On the **CORE / ACCESS PORT CONFIGURATION** tile, complete the port configuration for ESI-LAG between the collapsed core and access switches. The settings include a name and other port configuration elements. By default, this configuration includes the networks added on the **NETWORKS** tile on the same page. If you want to remove or modify the settings, click **Show Advanced** and configure the settings. Use the tips on the screen to configure the port profile settings.
- e. On the **DHCP RELAY** tile, configure the DHCP relay settings. You have the following options:
- **Enabled**—Configures DHCP relay on all the IRB-enabled devices in campus fabric. This option allows you to enable DHCP Relay on networks that you selected. The network will be populated inside the DHCP Relay tile as long as it is listed on the Networks tab on the same page.
  - **Disabled**—Disable DHCP relay on the devices in campus fabric. When you select this option, the DHCP relay is disabled on all the IRB-enabled devices. You should carefully select this option as this will remove the locally defined DHCP Relay on the Switch Detail page.
  - **None**—This option is automatically selected when the campus fabric topology has a mix of devices in terms of the DHCP relay configuration; that is, some devices have the DHCP relay enabled, some have it disabled, and some do not have it defined. This option will be visible for all Campus Fabric topologies that have DHCP Relay locally defined on individual switches.

If you want to remove all locally defined DHCP Relay networks, select **Enabled** and then choose **Remove all existing device level DHCP Networks**. You can simplify your DHCP Relay deployment by centralizing any configuration change from the campus fabric workflow.

If you enable DHCP relay in a campus fabric configuration, it is enabled on all the IRB-defined devices in the fabric and disabled on the rest of the devices. For example, in EVPN Multihoming topologies, DHCP relay is enabled on collapsed core devices and disabled on the rest.

10. Click **Continue** to go to the **Ports** tab, where you can configure the ports and create a connection between the core, distribution, and access layer switches.
11. Configure the switch ports in the collapsed core layer as follows:
  - a. Select a switch in the Collapsed Core section to open the switch port panel.
  - b. From the port panel of the switch, select a port that you want to configure.
  - c. Specify a port type (for example, ge or xe).
  - d. Select:
    - **Link to Collapsed Core** to connect the port to a core switch.
    - **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.



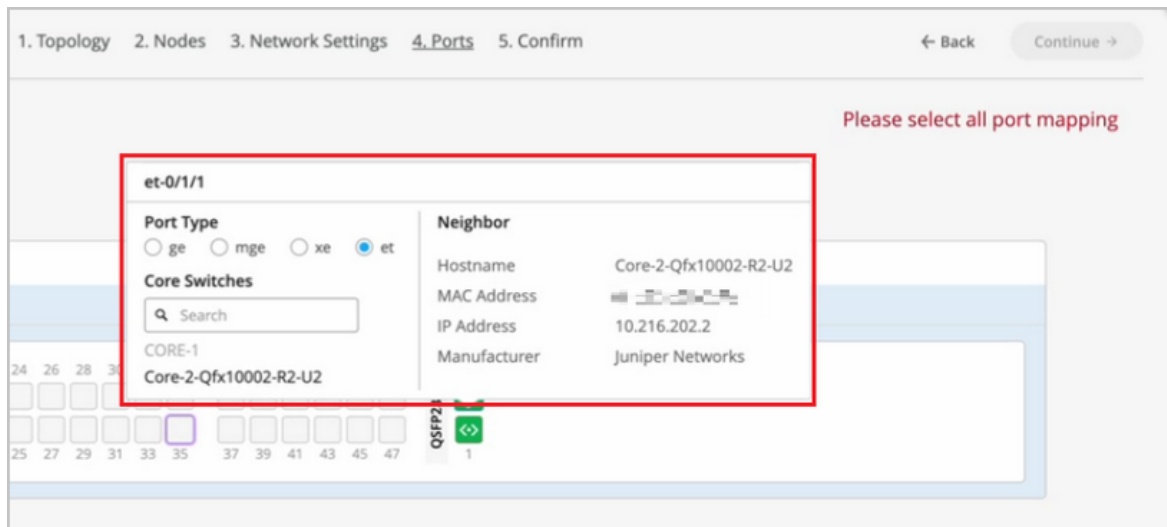
To configure the switch ports in the access layer:

- 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.



12. Click **Continue** to go to the **Confirmation** tab.
13. Click each switch icon to view and verify the configuration.
14. After verifying the configuration, click **Apply Changes > Confirm**.  
This step saves the campus fabric configuration to the Mist cloud and applies it to the switches. If the switches are offline, the configuration will be applied to them when they come online next time. A switch might take up to 10 minutes to complete the configuration.
15. Click **Close Campus Fabric Configuration**.  
After Mist builds the campus fabric, or while it is building the fabric, you can download the connection table. The connection table represents the physical layout of the campus fabric. You can use this table to validate all switch interconnects for the devices participating in the physical campus fabric build. Click **Connection Table** to download it (.csv format).

**BGP Summary**

Neighbor Information 2:43 PM (Updates Every 3 Minutes) 🔍

Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
Connected	Established	10.255.240.7	65003	65002	5m	5	2	21	17	default	Underlay
Connected	Established	192.168.255.21	65003	65002	5m	36	4	41	18	default	Overlay
Connected	Established	192.168.255.22	65004	65002	5m	36	40	42	39	default	Overlay

**Core1**

MAC Address f4:b5:2f:44:04:00  
 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 Status
Dist1-	1.2 GB	1.3 GB	Up
Dist2-	1.1 GB	1 GB	Up

Remote Shell Switch Insights

16. Verify the campus fabric configuration. To verify, follow the steps listed in the **Verification** section in [Campus Fabric EVPN Multihoming Workflow](#).



**Video:** [Deployment of Campus Fabric EVPN Multihoming With Wired Assurance](#)

## Configure Campus Fabric Core-Distribution

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.

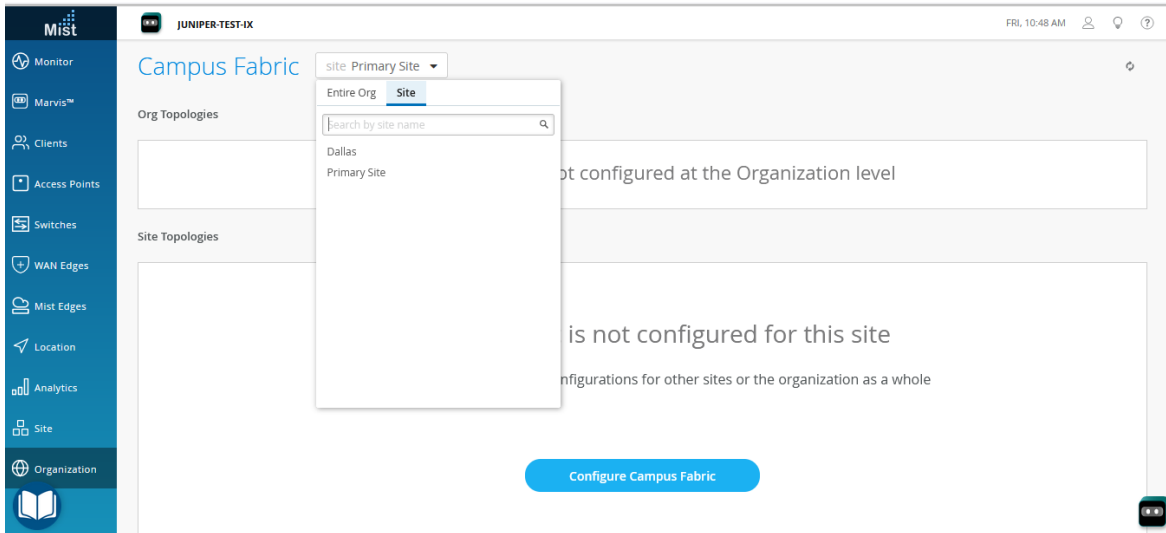
For more background information about campus fabric core-distribution architectures, see the following documents:

- [Campus Fabric Core Distribution CRB \(JVD\)](#)
- [Campus Fabric Core-Distribution ERB \(JVD\)](#)

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.





You can use an organization-level campus fabric topology to build a campus-wide architecture with multiple buildings. Otherwise, build a site-specific campus fabric with a single set of core, distribution, and access switches.

3. Click whichever option is relevant. Click the:
  - **Configure Campus Fabric** button (displayed if the site doesn't have a campus fabric configuration associated with it).
  - **Create Campus Fabric** button (displayed if the site already has at least one campus fabric configuration associated with it).


The **Topology** tab is displayed.

4. Select the topology type **Campus Fabric Core-Distribution**.


## Choose Campus Fabric Topology

Choose the topology you want to construct and configure related options


**TOPOLOGY TYPE**



**EVPN Multihoming**  
Collapsed core with ESI-Lag



**Campus Fabric Core-Distribution**  
EVPN core/distribution with ESI-Lag



**Campus Fabric IP Clos**  
Campus fabric with L3 at the edge

**CONFIGURATION**

Topology name is required

Topology Name

Topology Sub-type

CRB  
Centrally-routed and bridged with gateways on the Core

ERB  
Edge-routed and bridged with anycast gateways on the fabric edge

**TOPOLOGY SETTINGS**

BGP Local AS

(2-byte or 4-byte)

Underlay

IPv4  IPv6

Subnet ⓘ

(xxx::xxx/xx)

IPv6 Loopback Interface ⓘ

(xxx::xxx/xx)

IPv4 Auto Router ID Subnet / Loopback interface ⓘ

(xxx.xxx.xxx.xxx/xx)

Loopback Per-VRF Subnet ⓘ

(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 so that only the loopback IP addresses are exchanged in the underlay of the fabric.
  - **Underlay**—Select an internet protocol for the underlay. Options are IPv4 and IPv6. Only to ERB topologies support IPv6. You get the option to select IPv6 only if you selected ERB as Topology Sub-type.
  - **Subnet**— The range of IP addresses that Mist uses for point-to-point links between devices. You can use a range that suits your deployment. Mist breaks this subnet into /31 subnet addressing per link. You can modify this number to suit the specific deployment scale. For example, a /24 network would provide up to 128 point-to-point /31 subnets.
  - **IPv6 Loopback Interface**—Specify an IPv6 loopback interface subnet, which is used to autoconfigure IPv6 loopback interface on each device in the fabric.
  - **IPv4 Auto Router ID Subnet / Loopback Interface**—Mist uses this subnet to automatically assign a router ID to each device in the fabric (including access devices irrespective of whether they are configured with EVPN or not). Router IDs are loopback interfaces (lo0.0) used for overlay peering between devices. For new topologies, this field auto-populates a default subnet value (172.16.254.0/23), which you can modify. When you edit an existing

topology, this field doesn't populate any default value. The router ID is used as an identifier when deploying routing protocols such as BGP.

You can overwrite the automatically assigned router ID by manually configuring a loopback interface in the Router ID field on the Routing tile on the switch configuration page (**Switches** > **Switch Name**). However, if you modify the campus fabric configuration afterwards, Mist performs the automatic assignment of the router ID again, replacing the manually configured loopback interface.

- **Loopback per-VRF subnet**—Mist uses this subnet to automatically configure loopback interfaces (lo0.x) per the virtual routing and forwarding (VRF) instance that is used for services such as DHCP relay. For new topologies, this field auto-populates a default subnet value (172.16.192.0/24), which you can modify. This field supports a /19 or smaller subnet (for example, /24). When you edit an existing topology, this field doesn't populate any default value.
6. Click **Continue** to go to the **Nodes** tab, where you can select devices that form part of the campus fabric deployment.

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

To add switches:

- Click **Select Switches** in the section to which you want to add switches.

- b. Choose the switches that you want to add to the campus fabric.
- c. Click **Select**.

We recommend that you validate the presence of each device in the switch inventory before creating the campus fabric.

By default, Mist configures the core switches to function as border nodes that run the 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.

Also, Mist provides pods for improved scalability. Your access and distribution devices are grouped into pods. A pod could represent a building. For example, you can create a pod for each of the buildings in your site and create connections between the access and the distribution devices in that pod. You do not have to connect the same set of access devices to the distribution devices across multiple buildings. You can create multiple pods by clicking **+Add Nodes**.

You need only one connection between a pod and the core switch. You do not have to connect each distribution switch in a pod to all the core switches used. In a core-distribution topology (CRB or ERB), you need only one connection per core and distribution pair.

8. After selecting the switches, click **Continue** to go to the **Network Settings** tab, where you can configure the networks.
9. Configure the network settings, as described below:

- a. On the NETWORKS tile, add networks or VLANs to the configuration. You can either create a new network or import the network from the switch template defined on the **Organization > Switch** templates page.

To add a new VLAN, click **Create New Network** and configure the VLANs. The settings include a name, VLAN ID, and a subnet. You can specify IPv4 or IPv6 addresses for the subnet.

To import VLANs from the template:

- i. Click **Add Existing Network**.
- ii. Select a switch template from the **Template** drop-down list to view the VLANs available in that template.
- iii. Select the required VLAN from the displayed list, and click the ✓ mark.

VLANs are mapped to Virtual Network Identifiers (VNIs). You can optionally map 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. You can specify IPv4 or IPv6 addresses.
- d. On the DISTRIBUTION / ACCESS PORT CONFIGURATION tile, complete the port configuration for ESI-LAG between the collapsed core and access switches. The settings include a name and other port configuration elements. By default, this configuration includes the networks added on the NETWORKS tile on the same page. If you want to remove or modify the settings, click **Show Advanced** and configure the settings. Use the tips on the screen to configure the port profile settings.
  - e. On the DHCP RELAY tile, configure the DHCP relay settings. You have the following options:
    - **Enabled**—Configures DHCP relay on all the IRB-enabled devices in campus fabric. This option allows you to enable DHCP Relay on networks that you selected. The network will be

populated inside the DHCP Relay tile as long as it is listed on the Networks tab on the same page.

- **Disabled**—Disable DHCP relay on the devices in campus fabric. When you select this option, the DHCP relay is disabled on all the IRB-enabled devices. You should carefully select this option as this will remove the locally defined DHCP Relay on the Switch Detail page.
- **None**—This option is automatically selected when the campus fabric topology has a mix of devices in terms of the DHCP relay configuration; that is, some devices have the DHCP relay enabled, some have it disabled, and some do not have it defined. This option will be visible for all Campus Fabric topologies that have DHCP Relay locally defined on individual switches.

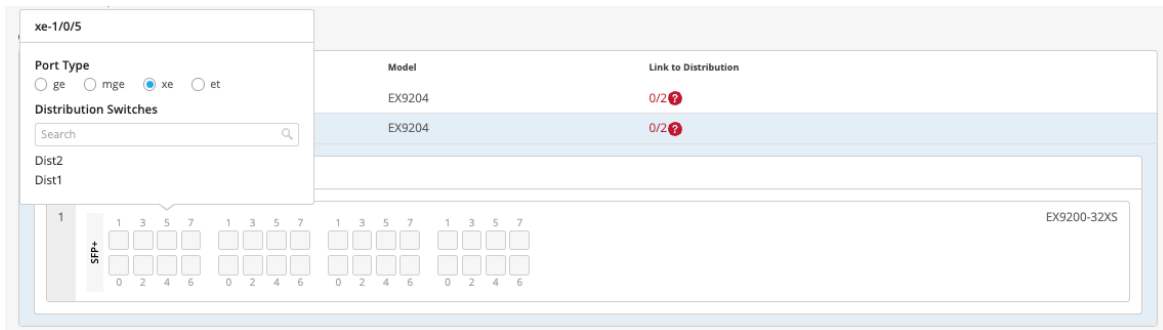
If you want to remove all locally defined DHCP Relay networks, select **Enabled** and then choose **Remove all existing device level DHCP Networks**. You can simplify your DHCP Relay deployment by centralizing any configuration change from the campus fabric workflow.

If you enable DHCP relay in a campus fabric configuration, it is enabled on all the IRB-defined devices in the fabric and disabled on the rest of the devices. For example, in Campus Fabric Core-Distribution (CRB) topologies, DHCP relay is enabled on core devices and disabled on the rest. Similarly, in Campus Fabric Core-Distribution (ERB), DHCP is enabled on distribution devices and disabled on the rest.

10. Click **Continue** to go to the **Ports** tab, where you can configure the ports and create a connection between the core, distribution, and access layer switches.

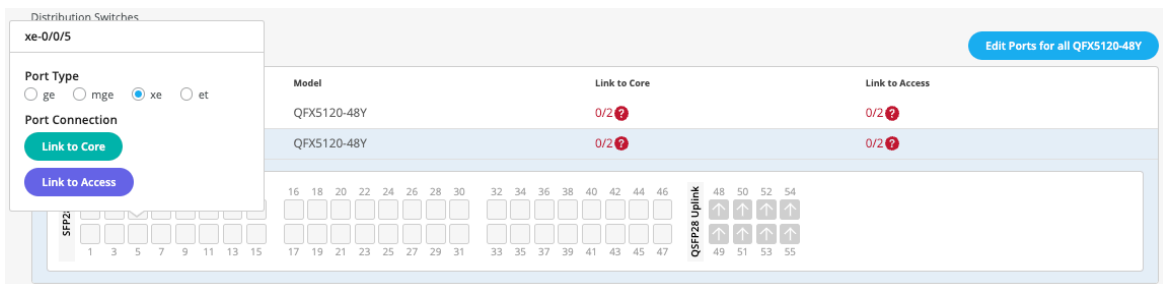
11. Configure the switch ports in the core layer as described below:
  - a. Select a switch in the Core section to open the switch port panel.
  - b. From the port panel of the core switch, select a port that you want to configure.
  - c. Specify a port type (for example, ge or xe).

- 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.



To configure switch ports in the distribution layer:

- Select a switch in the Distribution section to open the switch port panel.
- From the port panel of the switch, select a port that you want to configure.
- Specify a port type (for example, *ge* or *xe*).
- Select:
  - Link to Core** to connect the port to a core switch.
  - Link to Access** to connect the port to an access switch.
- 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.



To configure the switch ports in the access layer:

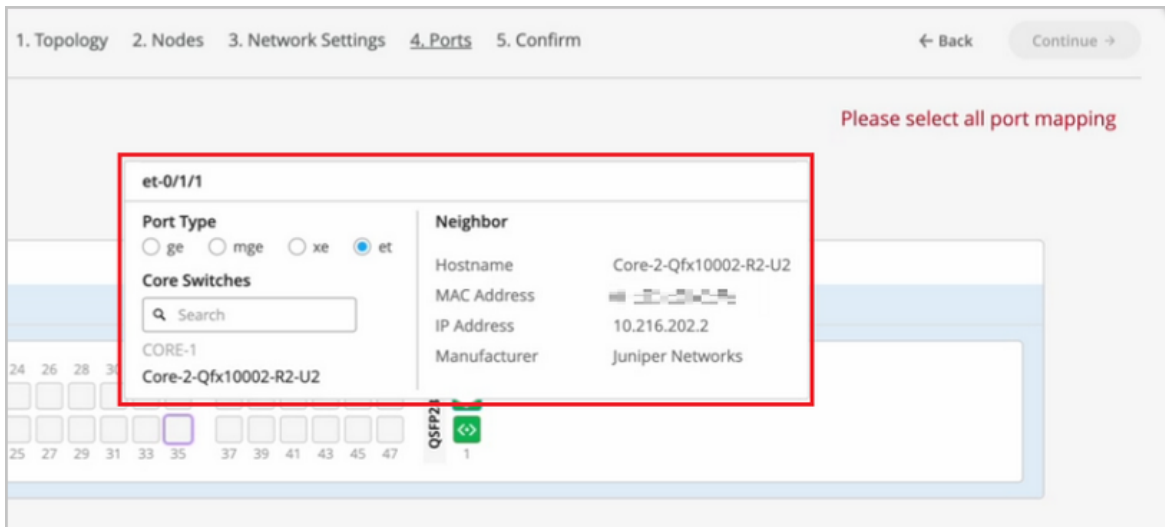
- Select a switch in the Access section to open the switch port panel.
- From the port panel of the switch, select a port that you want to configure.
- Specify a port type (for example, *ge* or *xe*).

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.



12. Click Continue to go to the **Confirmation** tab.
13. Click each switch icon to view and verify the configuration.
14. After verifying the configuration, click **Apply Changes > Confirm**.  
This step saves the campus fabric configuration to the Mist cloud and applies it to the switches. If the switches are offline, the configuration will be applied to them when they come online next time. A switch might take up to 10 minutes to complete the configuration.
15. Click **Close Campus Fabric Configuration**.  
After Mist builds the campus fabric, or while it is building the fabric, you can download the connection table. The connection table represents the physical layout of the campus fabric. You can use this table to validate all switch interconnects for the devices participating in the physical campus fabric build. Click **Connection Table** to download it (.csv format).

**BGP Summary**

Neighbor Information 2:43 PM (Updates Every 3 Minutes) 🔍

Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
Connected	Established	10.255.240.7	65003	65002	5m	5	2	21	17	default	Underlay
Connected	Established	192.168.255.21	65003	65002	5m	36	4	41	18	default	Overlay
Connected	Established	192.168.255.22	65004	65002	5m	36	40	42	39	default	Overlay

16. Verify the campus fabric configuration. To verify, follow the steps listed in the **Verification** section of [Campus Fabric Core Distribution CRB \(JVD\)](#) and [Campus Fabric Core-Distribution ERB \(JVD\)](#).

For a demo, watch the following video:



**Video:** [Deployment of Campus Fabric Core Distribution](#)



**Video:**

## 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.

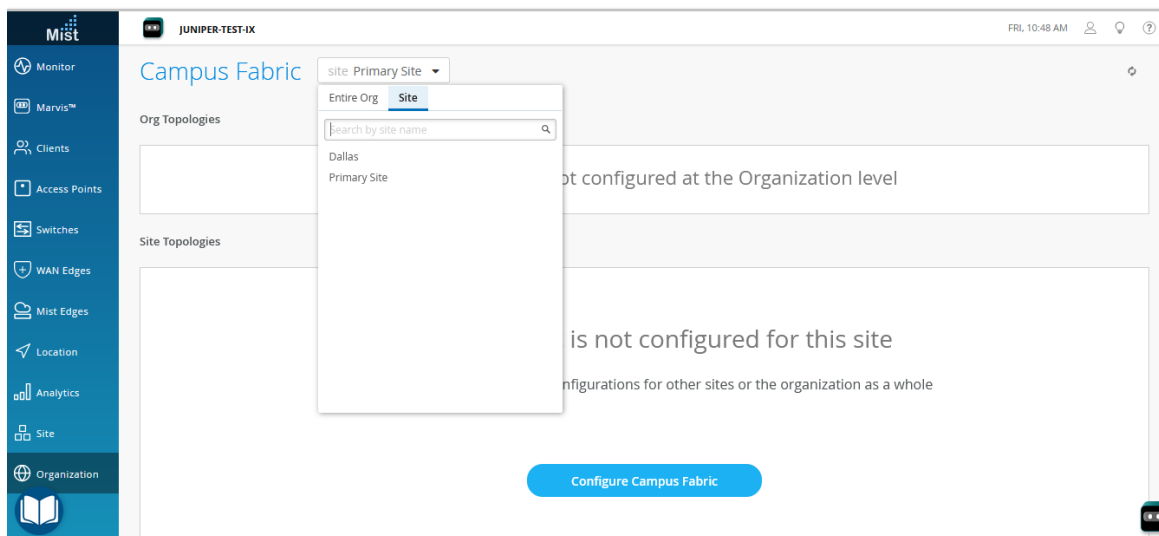
The campus fabric IP CLOS architecture supports Group Based Policies (GBPs) that enable you to achieve micro segmentation in the network. The GBP option gives you a practical way to create network access policies that are independent of the underlying network topology. In a GBP, you match a user group tag to a resource group tag to provide the specified users access to the specified resources. See ["Create a Switch Configuration Template"](#) on [page 31](#) to learn how to configure GBP on switches.

In a campus fabric IP Clos architecture, Mist provisions layer 3 (L3) integrated routing and bridging (IRB) interfaces on the access layer. All the access switches are configured with the same IP address for each L3 subnet. The end users terminating on the access layer have the default gateway set to the IRB address shared by all access layer devices. This deployment model utilizes anycast addressing for all devices participating in the L3 subnet. This deployment model provides a smaller blast radius for broadcast traffic and is ideal for east-west traffic patterns and IP Multicast environments.

For more detailed information about IP Clos architecture and its deployment, see [Campus Fabric IP Clos Using Mist Wired Assurance—Juniper Validated Design \(JVD\)](#).

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.



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 is required

Topology Name

**TOPOLOGY SETTINGS**

BGP Local AS

(2-byte or 4-byte)

**Underlay**

IPv4  IPv6

**Subnet** ⓘ

(xxx::xxx/xx)

**IPv6 Loopback Interface** ⓘ

(xxx::xxx/xx)

**IPv4 Auto Router ID Subnet / Loopback interface** ⓘ

(xxx.xxx.xxx.xxx/xx)

**Loopback Per-VRF Subnet** ⓘ

(xxx.xxx.xxx.xxx/xx)

5. Configure the topology name and other settings on the **Topology** tab, as described below:



**NOTE:** We recommend that you use the default settings on this screen unless they conflict with any networks attached to the campus fabric. The point-to-point links between each layer utilize /31 addressing to conserve addresses.

- a. In the **CONFIGURATION** section, enter the following:
- **Topology Name**—Enter a name for the topology.

b. (If you don't want to use the default settings) In the **TOPOLOGY SETTINGS** section, enter the following:

- **BGP Local AS**—Represents the starting point of private BGP AS numbers that are automatically allocated to each device. You can use any private BGP AS number range that suits your deployment. Mist provisions the routing policy so that only the loopback IP addresses are exchanged in the underlay of the fabric.
- **Underlay**—Select an internet protocol for the underlay. Options are IPv4 and IPv6.
- **Subnet**— The range of IP addresses that Mist uses for point-to-point links between devices. You can use a range that suits your deployment. Mist breaks this subnet into /31 subnet addressing per link. You can modify this number to suit the specific deployment scale. For example, a /24 network would provide up to 128 point-to-point /31 subnets.
- **IPv6 Loopback Interface**—Specify an IPv6 loopback interface subnet, which is used to autoconfigure IPv6 loopback interface on each device in the fabric.
- **IPv4 Auto Router ID Subnet / Loopback Interface**—Mist uses this subnet to automatically assign a router ID to each device in the fabric (including access devices irrespective of whether they are configured with EVPN or not). Router IDs are loopback interfaces (lo0.0) used for overlay peering between devices. For new topologies, this field auto-populates a default subnet value (172.16.254.0/23), which you can modify. When you edit an existing topology, this field doesn't populate a default value. The router ID is used as an identifier when deploying routing protocols such as BGP.

You can overwrite the automatically assigned router ID by manually configuring a loopback interface in the Router ID field on the Routing tile on the switch configuration page (**Switches > Switch Name**). However, if you modify the campus fabric configuration afterwards, Mist performs the automatic assignment of the router ID again, replacing the manually configured loopback interface.

- **Loopback per-VRF subnet**—Mist uses this subnet to automatically configure loopback interfaces (lo0.x) per virtual routing and forwarding (VRF) instance that is used for services such as DHCP relay. For new topologies, this field auto-populates a default subnet value (172.16.192.0/24), which you can modify. This field supports a /19 or smaller subnet (for example, /24). When you edit an existing topology, this field doesn't populate a default value.
6. Click **Continue** to go to the **Nodes** tab, where you can select devices that form part of the campus fabric IP Clos deployment.

**Campus Fabric Configuration** 1. Topology 2. Nodes 3. Network Settings 4. Ports 5. Confirm ← Back Continue →

**Select Campus Fabric Nodes** At least 1 access switch is required

Select the switches that will be used in each layer of the topology and provide Router IDs as required.

**Service Block Border**  Use Core as border ⓘ

**Core**

+  
Select Switches

**Pods** + Add Pod

Pod Name  
Pod 1

**Distribution**

+  
Select Switches

**Access**

+  
Select Switches

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

To add the switches:

- a. Click **Select Switches** in the section to which you want to add switches.
- b. Select the switches that you want to add to the campus fabric.
- c. Click **Select**.

We recommend that you validate the presence of each device in the switch inventory before creating the campus fabric.

By default, Mist configures the core switches to function as border nodes that run the 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. The minimum number of dedicated border nodes required is one.

Also, Mist provides pods for improved scalability. Your access and distribution devices are grouped into pods. A pod could represent a building. For example, you can create a pod for each of the

buildings in your site and create connections between the access and the distribution devices in that pod. You do not have to connect the same set of access devices to the distribution devices across multiple buildings. You can create multiple pods by clicking **+Add Nodes**.

You need only one connection between a pod and the core switch. You do not have to connect each distribution switch in a pod to all the core switches used. In an IPClos topology, you need only one connection between each core and distribution pair and between each distribution and access pair.

8. After selecting the switches, click **Continue** to go to the **Network Settings** tab, where you can configure the networks.
9. Configure the network settings, as described below.

- a. From the **NETWORKS** tile, add networks or VLANs to the configuration. You can either create a new network or import the network from the switch template defined in the **Organization > Switch** templates page.

To add a new VLAN, click **Create New Network** and configure the VLANs. The settings include a name, VLAN ID, and a subnet. You can specify IPv4 or IPv6 addresses for the subnet.

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:

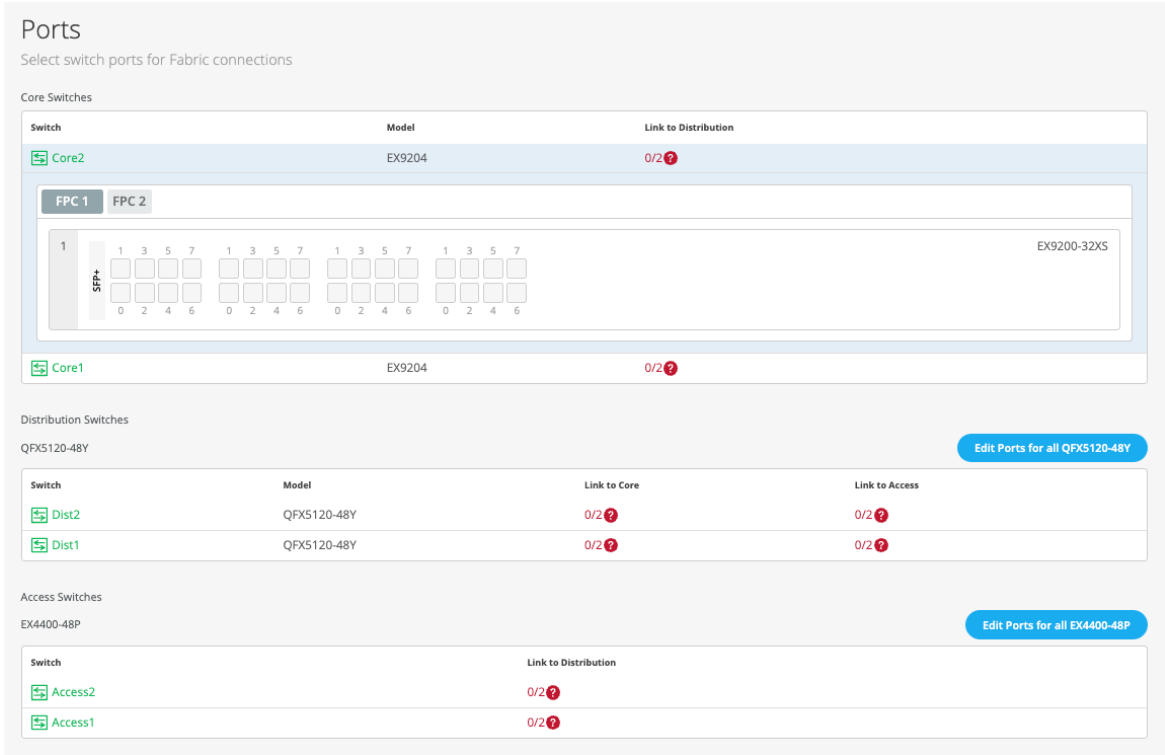
- i. Click **Add VRF Instance** and specify the settings. The settings include a name for the VRF and the networks to be associated with the VRF.
  - ii. To add extra routes, click the **Add Extra Routes** link on the **New VRF Instance** page and specify the route. You can specify IPv4 or IPv6 addresses.
- d. On the DHCP RELAY tile, configure the DHCP relay settings. You have the following options:
    - **Enabled**—Configures DHCP relay on all the IRB-enabled devices in campus fabric. This option allows you to enable DHCP Relay on networks that you selected. The network will be populated inside the DHCP Relay tile as long as it is listed on the Networks tab on the same page.
    - **Disabled**—Disable DHCP relay on the devices in campus fabric. When you select this option, the DHCP relay is disabled on all the IRB-enabled devices. You should carefully select this option as this will remove the locally defined DHCP Relay on the Switch Detail page.
    - **None**—This option is automatically selected when the campus fabric topology has a mix of devices in terms of the DHCP relay configuration; that is, some devices have the DHCP relay enabled, some have it disabled, and some do not have it defined.

If you want to remove all locally defined DHCP Relay networks, select **Enabled** and then choose **Remove all existing device level DHCP Networks**. You can simplify your DHCP Relay deployment by centralizing any configuration change from the campus fabric workflow.

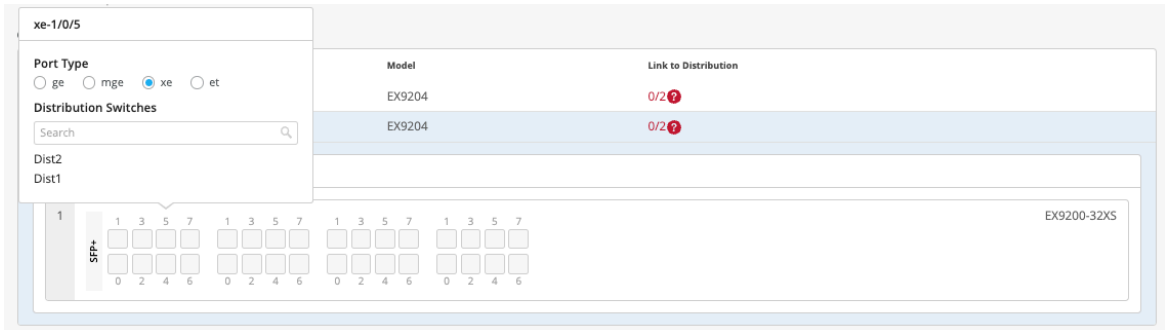
If you enable DHCP relay in a campus fabric configuration, it is enabled on all the IRB-defined devices in the fabric and disabled on the rest of the devices. For example, in Campus Fabric IP Clos edge topologies, DHCP is enabled on access devices and disabled on the rest.

10. Click **Continue** to go to the **Ports** tab, where you can configure the ports and create a connection between the core, distribution, and access layer switches.





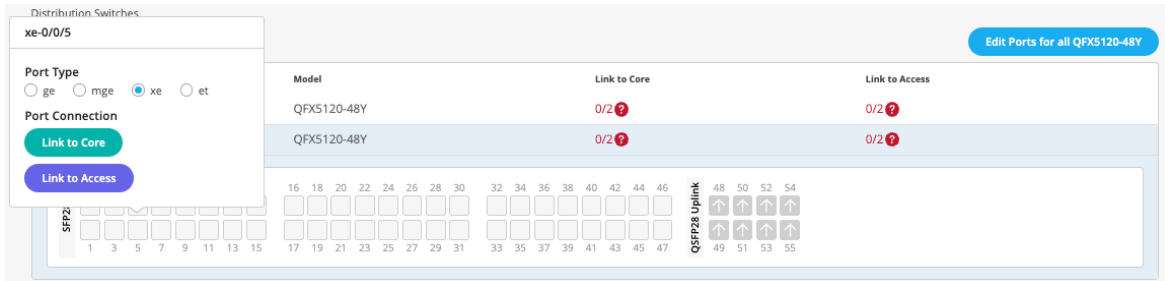
11. Configure the switch ports in the core layer as described below:
  - a. Select a switch in the Core section to open the switch port panel.
  - b. From the port panel of the core switch, select a port that you want to configure.
  - c. Specify a port type (for example, ge or xe).
  - d. Choose the distribution switch on which the link should terminate. You need to configure all the ports that need to be part of the campus fabric.



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.

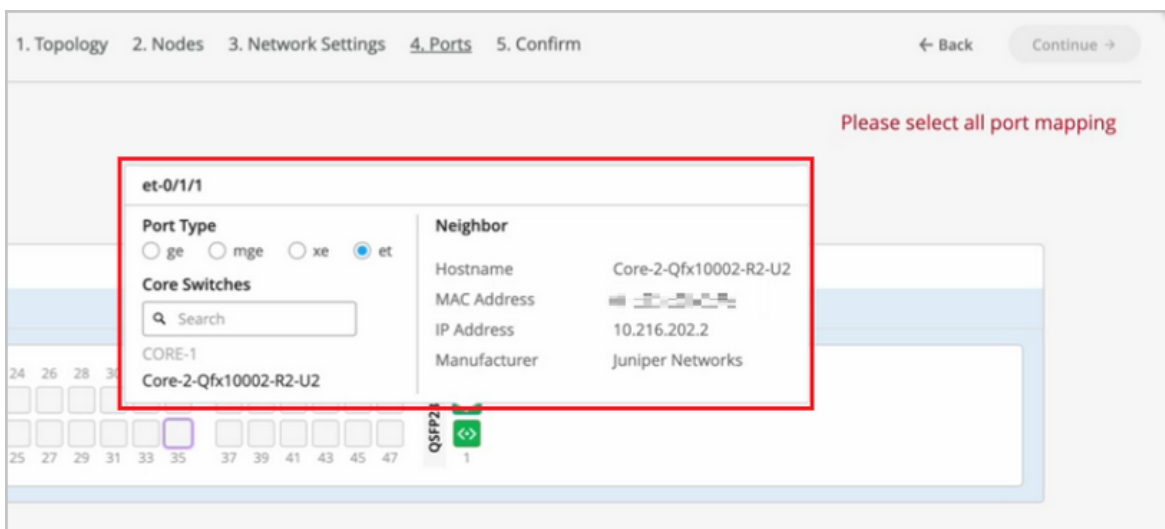
- c. Specify a port type (example: ge or xe).
- d. Select:
  - **Link to Core** to connect the port to a core switch.
  - **Link to Access** to connect the port to an access switch.
- e. Select the core or access switch (based on the selection in the previous step) on which the link should terminate. You need to configure all the ports that need to be part of the campus fabric.



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.



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).

**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

**Connections to Distribution**

Switch	RX Bytes	TX Bytes	Link Status
Dist1-	1.2 GB	1.3 GB	Up
Dist2-	1.1 GB	1 GB	Up

16. Verify the campus fabric configuration. To verify, follow the steps listed in the **Verification** section of [Campus Fabric IP CLOS Wired Assurance](#).

For a demo of the configuration steps, watch the following video:



**Video:** [Deployment of Campus Fabric IP CLOS](#)

After building an IP CLOS campus fabric, you can integrate it with a third party gateway (such as a router or firewall) by using BGP groups. Watch the following video for more information:



**Video:** [Integrate IP CLOS Fabric Using BGP](#)



CHAPTER

## Wired Service Levels

---

[Service Level Expectations \(SLE\) | 242](#)

[Wired SLEs Dashboard | 250](#)

[Wired Throughput SLE | 251](#)

[Wired Successful Connect SLE | 254](#)

[Switch Health SLE | 255](#)

[Switch Bandwidth SLE | 257](#)

---

# Service Level Expectations (SLE)

## SUMMARY

Get familiar with the Service Level Expectations (SLEs) and the SLE dashboard.

## IN THIS SECTION

- [What Are Service Level Expectations \(SLEs\)? | 242](#)
- [Finding the SLE Dashboard | 243](#)
- [Selecting the Context and Time Period | 243](#)
- [Using the System Changes Timeline | 245](#)
- [Setting the SLE Thresholds | 246](#)
- [Understanding the SLE Blocks | 247](#)
- [Sample SLE Block | 248](#)
- [Viewing the Root Cause Analysis Page | 249](#)

## What Are Service Level Expectations (SLEs)?

The following video gives you a quick, high-level introduction to SLEs.



**Video:** [Mike and Marvis Episode 2: Understanding Service Level Expectations](#)

Juniper Mist™ captures, analyzes, correlates, and classifies event and performance data from your network and devices. It then provides you with an assessment of the quality of users' experiences on your network.

Many factors contribute to positive or negative user experiences. Juniper Mist organizes these factors into Service Level Expectations (SLEs). You can set the SLE thresholds to define exactly what "success" means for SLEs such as throughput, capacity, AP health, switch health, and more (as relevant to your network).

When user experiences fail to meet your SLE success thresholds, Juniper Mist identifies the root cause of each poor experience and provides complete details so that you can address the issues.

By skimming the SLE dashboard, you can see at a glance which service levels are low and what types of issues need to be addressed.

## Finding the SLE Dashboard


To access an SLE dashboard, select **Monitor** > **Service Levels** from the left menu. Then use the buttons at the top of the page to select the dashboard that you want to view (such as Wireless, Wired, WAN, Location, and Insights).



**NOTE:** Your subscriptions determine which buttons appear (for example, you need a Juniper Mist Wi-Fi Assurance subscription for Wireless SLEs).

## Selecting the Context and Time Period

At the top of the Monitor page, select the context, which can be an entire organization, an access point, or a client. In addition, select a time period, such the last 60 minutes, the last 7 days, or a date range.



**NOTE:** The Monitor page displays data as recent as the past 60 minutes or as far back as the last 7 days. If you purchase a Premium Analytics subscription, you can access up to 3 years' worth of wireless network insights and other data. To access the information available through your Premium Analytics subscription, select **Analytics** > **Premium Analytics** from the left menu of the portal.

### Context Example: Organization

To compare the performance of all sites in your organization, select **Entire Org** as the context.

Site	Avg AP Count	Avg Client Count	Overall Service	Time to Connect	Successful Connect	Coverage	Roaming	Throughput	Capacity	AP Health
Live-Demo	15	11	85%	96%	62%	93%	98%	100%	55%	92%
Westford	1	1	> 99%	100%	100%	100%	--	100%	> 99%	100%

(includes up to 100 sites, excludes sites with no data for the selected Service Level)

Use the filter buttons above the table to change the view:

- Overall Service—This is the default view when you select Entire Org as the context. You can compare the overall user experience at each site.
- SLE filter buttons—Zoom in on a single SLE by using the SLE buttons above the table. The button options vary, depending on which page you're on (Wireless, WAN, and so on).

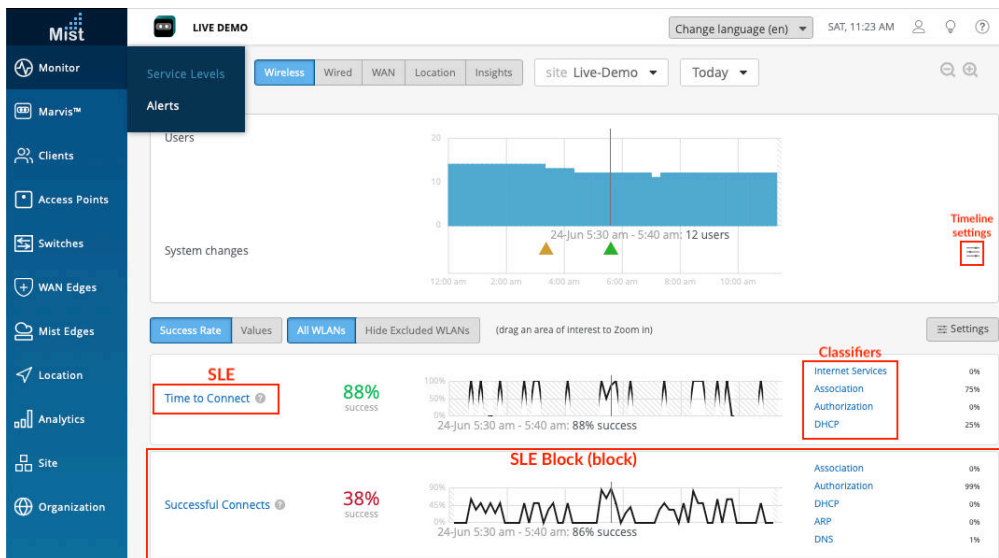
You also have the option to view **All Sites** or the **Worst 100 Sites**. For the Worst 100 option, also use the drop-down list to select the SLE that you're concerned about. For example, if you're troubleshooting an issue with capacity, you'd select that option from the drop-down list to see which sites are having the most issues with this SLE.



**NOTE:** The available SLEs for the filter buttons and the Worst 100 drop-down list vary, depending on whether you're looking at Wireless, Wired, or WAN SLEs.

### Context Example: Site


To compare all SLEs for one site, select the site as the context.



**NOTE:** This image shows a Wireless example, but the SLE blocks are set up the same way for Wired, WAN, and so on.

## Using the System Changes Timeline

When investigating issues, your first question might be, "Did anything change on the network?" With this timeline, you can see at a glance if any system changes occurred and how many users or clients were active at the time.



The triangles below the timeline represent various types of system changes:

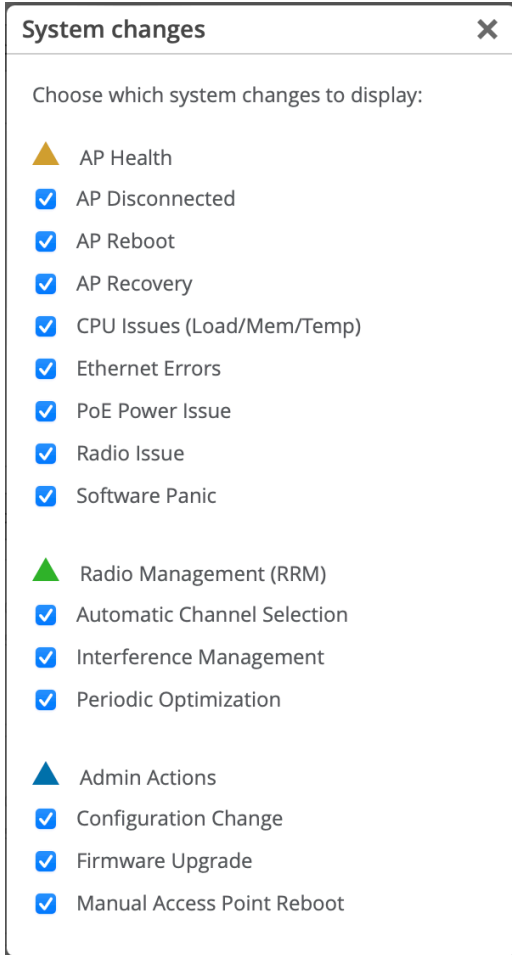
- Yellow triangle—AP Health
- Green triangle—Radio Management (RRM)
- Blue triangle—Admin Actions

You can adjust the timeline settings to specify the types of changes to include. To get started, click the timeline settings button:



In the System Changes window, select or deselect check boxes for each event that you want to include or exclude.





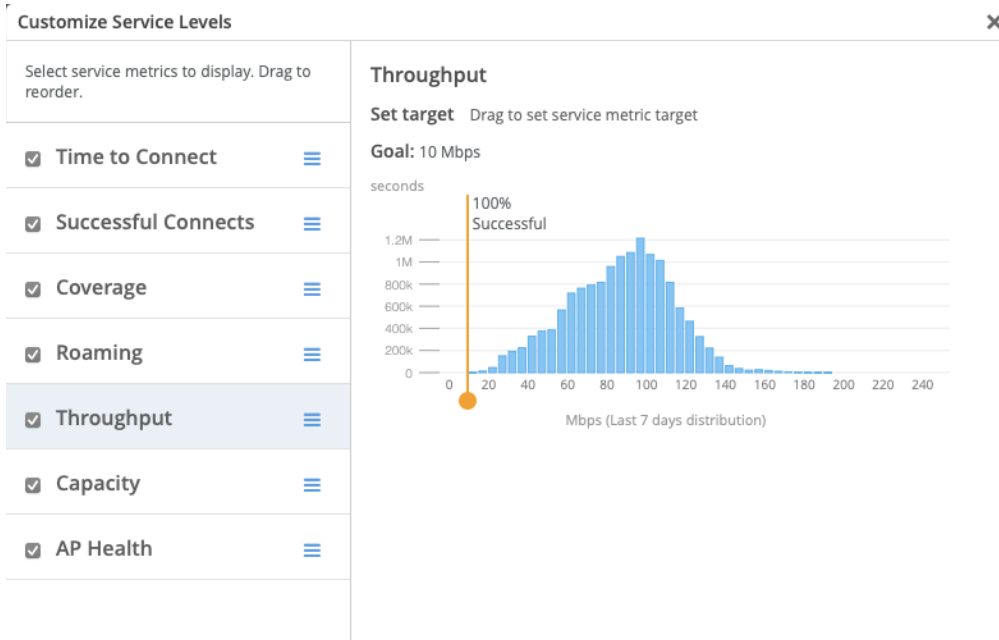
## Setting the SLE Thresholds

Each SLE has a success threshold. For the Time to Connect SLE, for example, you might set a threshold of 2 seconds. This means that you consider your network successful when users can send and receive data over the Internet within 2 seconds of attempting to associate with an access point.

To view or modify the SLE thresholds, you can click the **Settings** button on the right side of the SLE dashboard.



In the Customize Service Levels window, you can modify the thresholds as needed to ensure that the SLE settings meet your goals for your network.



**NOTE:** This example shows the wireless SLEs. Depending on the dashboard that you're viewing, you'll see different SLEs in this window.

## Understanding the SLE Blocks

Each SLE is represented by a separate block (sub-section) on the dashboard.

In each block, you'll see:

- **Overall Service Level.** On the left side of each SLE block, you'll see the overall service level for the selected site and time period.
  - Click **Success Rate** to see the *percentage* of user experiences that met the SLE success threshold.
  - Click **Values** to see the *number* of user experiences that met the SLE success threshold.
- **Timeline.** In the middle of each SLE block, you can explore the timeline. As your mouse moves across the timeline, information appears under it.
  - Click **Success Rate** to see the *percentage* of successful user experiences at the selected point in time.
  - Click **Values** to see the *number* of successful user experiences at the selected point in time.

- **Classifiers.** On the right side of each SLE block, you see the *classifiers* for the user experiences that didn't meet the SLE success threshold. Juniper Mist attributes each unsuccessful user experience to one classifier. Together, the classifiers give you a high-level root cause analysis of the unsuccessful user experiences.
- Click **Success Rate** to see the *percentage* of unsuccessful user experiences that were caused by each classifier.




**NOTE:** Together, these individual percentages total 100 percent of the unsuccessful user experiences.

- Click **Values** to see the *number* of unsuccessful user experiences that were caused by each classifier.



**NOTE:** Together, these individual values represent the total number of unsuccessful user experiences.

## Sample SLE Block



In this example, the Success Rate button is selected, so you see percentages instead of values.

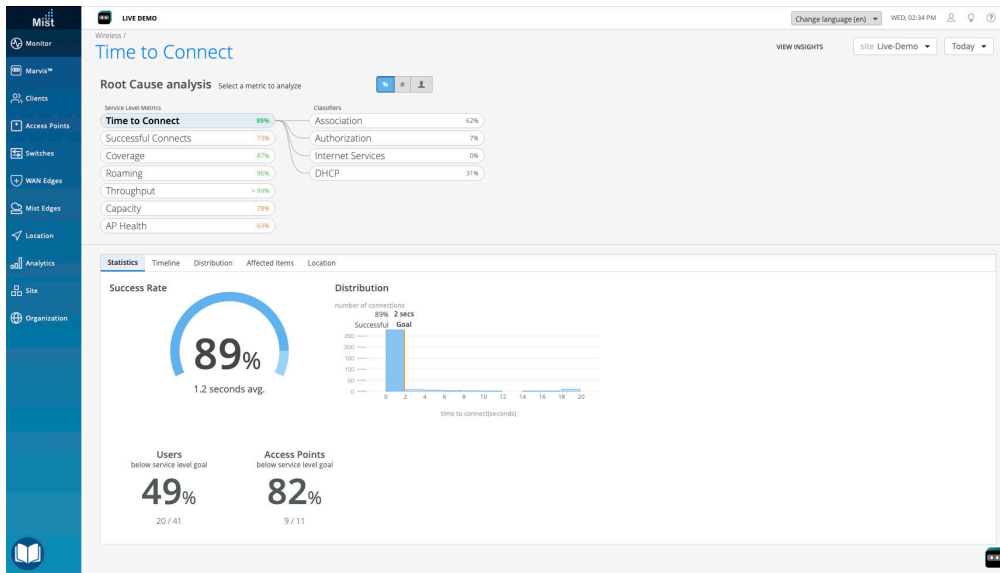
- On the left, you see that the overall success rate for the selected site and time period was 88 percent.
- In the middle, the timeline caption shows that the mouse is hovering over 24-Jun 5:30 am - 5:40 am. At that point, the success rate was 88 percent.

On the right, you see that 75 percent of the SLE-lowering issues occurred in the Association process and 25 percent occurred in the DHCP process. Together, these classifiers account for 100 percent of the user experiences that failed to meet the threshold. The other classifiers show 0 percent, meaning that they did not have any impact on this SLE.

## Viewing the Root Cause Analysis Page

From the dashboard, you can click any SLE or classifier to go to the Root Cause Analysis page.

This example shows the Root Cause Analysis page for the wireless Time to Connect SLE.



### Tips:

- At the top of the page, you see the data for all classifiers and their sub-classifiers (if applicable).
- In the lower part of the page, you see additional details about the selected item. Depending on the classifier, you might see signal strength information, a list of affected devices and clients, or other information. These details help you to understand the scope of the issues.
- On the Affected Items page, you can use the Filter box to search for an item. As shown in the animation below, simply start typing in the box, and matching items will appear in a drop-down list. Then click the item that you want to view.

The screenshot shows a dashboard with tabs for Statistics, Timeline, Distribution, **Affected Items**, and Location. Under 'Affected Items', there are two summary cards: 'Users' with a count of 9 and 'Access Points' with a count of 6. A search bar labeled 'Filter' is present. Below is a table of affected items:

Name	Overall Impact	Failure Rate	MAC Address
hal	10.00%	100%	dca6:32:c7:e7:e6
denali	20.00%	100%	50:32:37:ea:c3:c2
ac:67:84:0e:d4:74	10.00%	100%	ac:67:84:0e:d4:74
abhiramms-mbp	10.00%	50%	88:66:5a:18:2d:1f
prajendir-P16	10.00%	50%	30:89:4a:df:ec:6f
satishj-mbp	10.00%	33%	bc:d0:74:59:bd:c2
svadi-mbpm1	10.00%	33%	bc:d0:74:15:82:54
rdandamudi-mbp	10.00%	25%	bc:d0:74:7e:14:7a
fe:29:6e:cc:16:ac	10.00%	13%	fe:29:6e:cc:16:ac

## Wired SLEs Dashboard

### SUMMARY

Get started using the wired service-level experience (SLE) dashboard to assess the service levels for user-impacting factors such as throughput, connectivity, and switch health.


### IN THIS SECTION

- [Finding the Wired SLEs Dashboard | 251](#)
- [Wired Assurance: Day 2 - Wired Service Level Expectations \(SLEs\) Video Overview | 251](#)
- [Using the Wired SLE Dashboard | 251](#)

Juniper Mist™ 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.

## Finding the Wired SLEs Dashboard

To find the Wired SLEs dashboard, select **Monitor** > **Service Levels** from the left menu, and then click the **Wired** button.



**NOTE:** The buttons appear only if you have the required subscriptions. For information about these requirements, see the [Juniper Mist AI-Native Operations Guide](#).

## Wired Assurance: Day 2 - Wired Service Level Expectations (SLEs) Video Overview



Video: [Wired Assurance: Day 2 - Wired Service Level Expectations \(SLEs\)](#)

## Using the Wired SLE Dashboard

For a general introduction to SLEs, see "[Service Level Expectations \(SLE\)](#)" on page 242.

For help interpreting the wired SLEs and classifiers, explore the other Wired SLE topics in this chapter.

# Wired Throughput SLE

### SUMMARY

Use the Wired Throughput SLE to assess users' experiences with throughput on your wired network.

### IN THIS SECTION

- [What Does the Wired Throughput SLE Measure? | 252](#)

Throughput is one of the Service-Level Expectations (SLEs) that you can track on the wired SLEs dashboard.



**NOTE:** To find the Wired SLEs dashboard, select **Monitor** > **Service Levels** from the left menu, and then click the **Wired** button.

## What Does the Wired Throughput SLE Measure?

This SLE represents the ability of wired users to pass traffic without impedance.

## Classifiers

When the throughput threshold is not met, Juniper Mist sorts the issues into classifiers. The classifiers appear on the right side of the SLE block. In this example, less than 1 percent of the issue were attributed to Congestion Uplink, 19 percent to Interface Anomalies, 1 percent to Storm Control, and 80 percent to Congestion. (See the classifier descriptions below the example.)



- **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 a Spanning Tree Protocol (STP) root port.
  - The uplink port has a higher number of transmitted and received packets compared to the other ports.
  - Aggregated Links. Congestion can also be caused by aggregated Ethernet links and module ports.

- **Interface Anomalies**—The details for interface anomalies are all obtained from the switch. The Interface Anomalies classifier contains three sub-classifiers: MTU Mismatch, Cable Issues, and Negotiation Failed.
  - **MTU Mismatch**—As an administrator, you can set an MTU value for each interface. The default value for Gigabit Ethernet interfaces is 1514 . To support jumbo frames, you must configure an MTU value of 9216, which is the upper limit for jumbo frames on a routed virtual LAN (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 Networks 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 autonegotiation failure, duplex conflicts, or user misconfiguration of device settings. Moreover, older devices may fail 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 level. This threshold is known as a storm control level or storm control bandwidth. The default storm control level is 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.
- **Congestion**—This classifier measures the number of output drops. When packets come into a switch interface, they are placed in an input queue (buffer). When the buffer becomes full, it will start to drop packets (TxDrops). We use a formula that takes into account the following ratios to determine if there is a 'bad user minute' due to 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.



# Wired Successful Connect SLE

## SUMMARY

Use the Wired Successful Connect SLE to assess clients' experiences connecting to your wired network.

## IN THIS SECTION

- [What Does the Wired Successful Connect SLE Measure? | 254](#)
- [Classifiers | 254](#)

Successful Connect is one of the Service-Level Expectations (SLEs) that you can track on the Wired SLEs dashboard.



**NOTE:** To find the Wired SLEs dashboard, select **Monitor** > **Service Levels** from the left menu, and then click the **Wired** button.

## What Does the Wired Successful Connect SLE Measure?

Juniper Mist monitors client connection attempts and identifies failures. This SLE helps you to assess the impact of these failures and to identify the issues to address.



**NOTE:** This SLE will show data only if you use 802.1X on the wired network to authenticate clients or if you have DHCP snooping configured.

## Classifiers

When connection attempts are unsuccessful, Juniper Mist sorts the issues into classifiers. The classifiers appear on the right side of the SLE block. In this example, 100 percent of the issues are attributed to Authentication. (See the classifier descriptions below the example.)



- **DHCP**—Dynamic Host Configuration Protocol (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 (DCHP timeouts).



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

- **Authentication**—Each time a client authenticates, a client event is generated. These could either be successful or failed events. This classifier helps you identify issues that caused authentication failures. Here's a list of possible reasons for an 802.1X 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.
    - A routing issue exists between the switch and the RADIUS server.
    - The RADIUS server is down.
  - If all switch ports on all the 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 a Windows device, fails to authenticate, it may suggest an issue related to certifications.

## Switch Health SLE


### SUMMARY

Use the Switch Health SLE to assess switch performance and to identify user-impacting issues with switch reachability, memory, CPU, and more.

### IN THIS SECTION

- [What Does the Switch Health SLE Measure? | 256](#)

Switch Health is one of the Service-Level Expectations (SLEs) that you can track on the Wired SLEs dashboard.




**NOTE:** To find the Wired SLEs dashboard, select **Monitor** > **Service Levels** from the left menu of the Juniper Mist™ portal, and then select the **Wired** button.

## What Does the Switch Health SLE Measure?

Juniper Mist™ monitors your switches' operating temperatures, power consumption, CPU, and memory usage. Monitoring switch health is crucial because issues such as 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.

## Classifiers

When the Switch Health threshold is not met, Juniper Mist sorts the issues into classifiers. The classifiers appear on the right side of the SLE block. In this example, 82 percent of the issues are attributed to Switch Unreachable and 12 percent to System. (See the classifier descriptions below the example.)



- **Switch Unreachable**—The switch can't be accessed.
- **Capacity**
  - **ARP Table**—Usage exceeded 80 percent of the Address Resolution Protocol (ARP) table capacity.
  - **Route Table**—Usage exceeded 80 percent of the routing table capacity.


- **MAC Table**—Usage exceeded 80 percent of the MAC table capacity.
- **Network**—You can use this classifier to monitor user minutes when the throughput is lower than expected due to uplink capacity limitations. 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 identify these issues:
  - **WAN Latency**—Displays user minutes affected by latency. The latency value is calculated based on the average value of RTT over a period of time.
  - **WAN 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.
- **System**
  - **CPU**—The CPU usage of the switch is above 90 percent.
  - **Memory**—The memory utilization is above 80 percent.
  - **Temp**—The operating temperature of the switch is outside the prescribed threshold range, going either above the maximum limit or below the minimum requirement.
  - **Power**—The switch is consuming over 90 percent of the available power.

## Switch Bandwidth SLE

### IN THIS SECTION

- [What Does the Switch Bandwidth SLE Measure? | 258](#)
- [Classifiers | 258](#)

Switch Bandwidth is one of the Service-Level Expectations (SLEs) that you can track on the Wired SLEs dashboard.



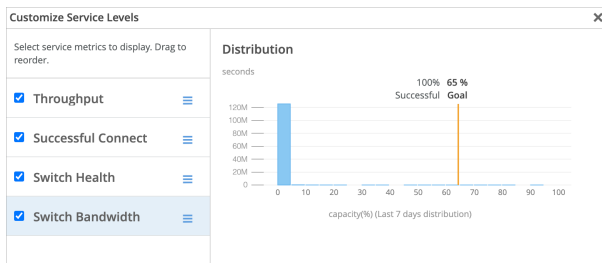


**NOTE:** To find the Wired SLEs dashboard, select **Monitor** > **Service Levels** from the left menu, and then click the **Wired** button.

## What Does the Switch Bandwidth SLE Measure?


Juniper Mist™ measures the available bandwidth on your network based on the queued packets and dropped packets for each configured queue. The ratio between total\_DroppedPackets and total\_QueuedPackets is used to determine congestion at the interface level. The most dropped queue is also noted in the details for distribution/affected items. This SLE can help you to determine if you need more wired bandwidth on your site.

You can click the **Settings** button (above the SLE blocks) to set the percentage to use as the success threshold for this SLE. The percentage represents the total\_DroppedPackets as a portion of total\_QueuedPackets.



## Classifiers

When the Switch Bandwidth threshold is not met, Juniper Mist sorts the issues into classifiers. The classifiers appear on the right side of the SLE block. In this example, 33 percent of the issues are attributed to Congestion and 67% to Bandwidth Headroom. (See the classifier descriptions below the example.)



- **Congestion**—This classifier measures the number of output drops. When packets come into a switch interface, they are placed in an input queue (buffer). When the buffer becomes full, it will start to drop packets (Tx Drops). We use a formula that takes into account the following ratios to determine if there are bad user minutes due to 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.
- **Bandwidth Headroom**—This classifier is triggered if the bandwidth usage exceeds the threshold for this SLE.
- **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 a Spanning Tree Protocol (STP) root port.
  - The uplink port has a higher number of transmitted and received packets compared to the other ports.
  - There is congestion due to aggregated Ethernet links and module ports.

# 7

CHAPTER

## Troubleshooting

---

[Troubleshoot Your Switch Connectivity](#) | 261

[Troubleshooting Juniper CloudX](#) | 267

[Cloud-Ready LED Blink Patterns](#) | 271

[Troubleshoot with Marvis](#) | 275

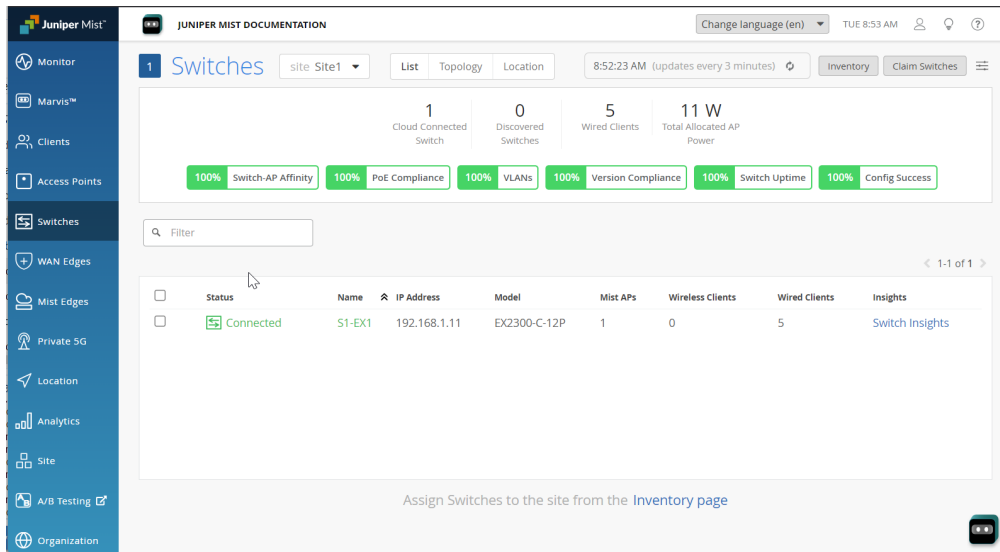
[FAQs \(Mist Wired\)](#) | 276

---

# Troubleshoot Your Switch Connectivity

You can get real-time statistics for a given switch interface by selecting the switch and clicking the **Live Traffic Counters** button, or drilling down on a specific interface or virtual chassis member and then doing the same. Traffic statistics include input and output, L2 and L3 errors, and BUM traffic.

Figure 23: Live Traffic Counters

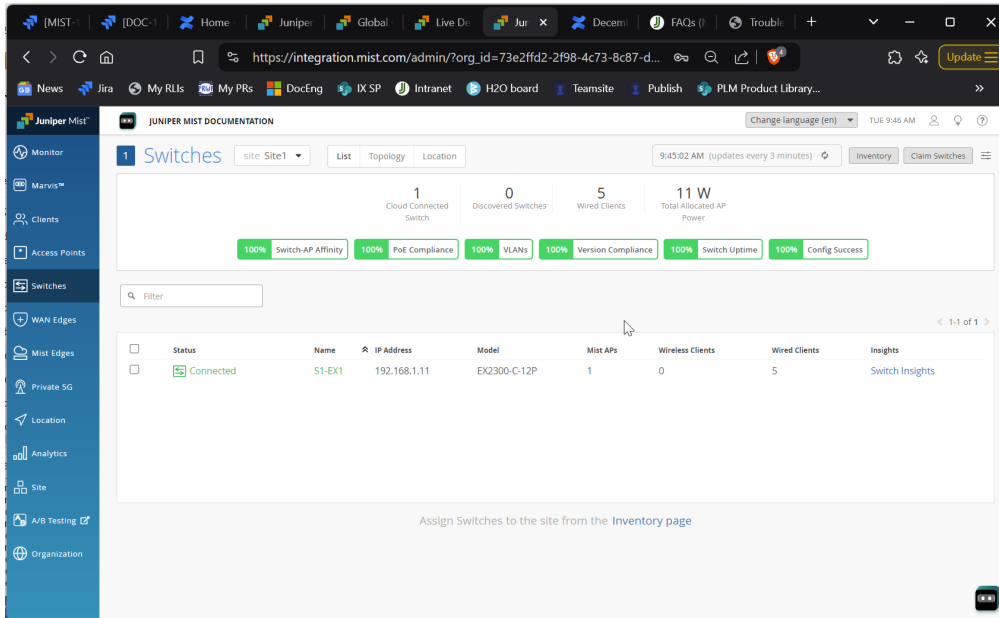


## Viewing Switch Processes

Real-time statistics for the processes running on a given switch are available by selecting the switch and then drilling down to the switch Insights page. Click the **View Live Process Detail** button. Supports virtual chassis and virtual machines in addition to physical devices.



Figure 24: Live Process Detail



## Troubleshooting Switches

If the Juniper Mist™ portal shows a switch as disconnected when it is online and reachable locally, you can troubleshoot the issue. You need console access or SSH access to the switch to perform the troubleshooting steps listed in this topic.

To troubleshoot your switch:

1. Ensure that the Junos OS version running on the switch supports zero-touch provisioning (ZTP). For example, the EX2300 and EX3400 switches require Junos OS version 18.2R3-S2 or later. The EX4300 switch requires Junos OS 18.4R2-S2 or later. The EX4600 and EX4650 switches require Junos OS 20.4R3 or later.
2. Log in to the switch CLI and run `show interfaces terse`.

```
user@switch> show interface terse
Interface      Admin Link Proto      Local
ge-0/0/0      up    up
irb.0         up    up    inet      192.168.3.24/24
me0          up    down
me0.0        up    down  inet      192.168.3.24/24
...truncated...
```

You should see the integrated routing and bridging (IRB) interface (irb.0) with an IP address. You might see multiple IRB interfaces, depending on the switch model (or in the case of a Virtual Chassis).

At least one IRB interface needs to have a valid IP address. The switch can also connect using a management IP address, which you can see on the me0 interface. Ensure that either the irb0 or me0 interface has a valid IP address and has its Admin and Link states up.

3. Ensure that the switch can reach the gateway.
4. Use a ping test, as follows, to ensure that the switch can reach the Internet:

```
user@switch> ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8): 56 data bytes
64 bytes from 8.8.8.8: icmp_seq=0 ttl=117 time=22.996 ms
64 bytes from 8.8.8.8: icmp_seq=1 ttl=117 time=24.747 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=117 time=16.528 ms

--- 8.8.8.8 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 16.528/21.424/24.747/3.535 ms
```

5. Check if the switch can resolve oc-term.mistsys.net and jma-term.xx.mistsys.net by using a ping test. Sample ping tests are shown below:

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

```
user@switch> ping jma-terminator-staging.mistsys.net
PING a8481a00030ad459aac15af07d5f2c5b-75855524.us-east-1.elb.amazonaws.com (3.210.247.53):
56 data bytes
^C
--- a8481a00030ad459aac15af07d5f2c5b-75855524.us-east-1.elb.amazonaws.com ping statistics
---
1 packets transmitted, 0 packets received, 100% packet loss
```

If the switch is not resolving `oc-term.mistsys.net` or `jma-term.xx.mistsys.net`, make sure that the switch has a DNS server configured.

```
user@switch> show configuration | display set | grep name-server
set system name-server 202.56.230.2
set system name-server 202.56.230.7
set system name-server 8.8.8.8
```

If the switch doesn't have a DNS server, configure the server as shown in the following example:

```
user@switch# set system name-server 8.8.8.8
```

6. Ensure that the required firewall port (TCP port 2200 for `oc-term.mistsys.net`) is open.

```
user@switch> show system connections | grep 2200
tcp4 0 0 192.168.3.24.64647 13.56.90.212.2200 ESTABLISHED
```

See [Device-to-Cloud Addresses and Ports](#) to determine which port to enable, depending on your cloud environment.



**NOTE:** The EX2300, EX3400, EX4100, EX4400, EX4650, EX5120 switches no longer need the port 2200. These switches connect to Mist cloud over HTTPS port 443. See also: "[Troubleshooting Juniper CloudX](#)" on page 267.

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-
ws5Qz6tuRhSv8Xr187dVY2TzF/u0EcyKWlleUjikPfIEhSrvxNdbYgRhK8x7Vbk.mf5F9Cu0BEtp0IcSMWoJZjmfFn/
CA05TIEhSeK4aJUjqP5Q9tu4an/Ct0B7-dboJZUjHmfaJn/ApREevW8X-
YgoiqmxNb2gaUD69Cp1RSyKMLxCt0RSrvM7-VboJDjqPTzNdmfzF/
9vW8LdbY2aZGisY4ZDif5z3690By1KWX7KvZUHkTQ1KvW-VJGDiqmGU/
CtuEhKM87wYaJdkqfoaQFn6At1RhrM8xNd"; ## SECRET-DATA
keep-alive {
retry 3;
timeout 5;
}
services netconf;
oc-term.mistsys.net {
port 2200;
retry 1000;
timeout 60;
}
}

```

See [outbound-ssh](#) for more information.

You can also examine the log messages by using the command `show log messages`.

9. If you are adding the switch for the first time, do the following:
- Delete the present Juniper Mist configuration from the switch using the delete command.
  - Onboard the switch again using the claim or adopt workflow.
  - Verify the system connection using the `show system connections | grep 2200` command. If the switch remains disconnected with the sessions stuck in FIN\_WAIT state, but is able to reach the Internet and resolve DNS, check for any maximum transmission unit (MTU) issues.
10. To check for any MTU issues, initiate a ping test toward any public server (for example, 8.8.8.8).

Another way to check for MTU issues is to review the uplink packet capture file from the switch. A failing transaction due to an MTU issue would look like the following example. The example shows that the packets with a size of 1514 are being retried.

Time	Source	SrcPort	Destination	DstPort	Protocol	Transaction ID	Length	TTL	TCP/STMP	Response time	Seq#	Len	RetSeq#	ACK#	BytDiff	CalcDiff	TCPdelta	Identification	Transaction ID	Your (client) IP address	ACK#
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	66680	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	66680	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	29671	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	489	64	8	296..	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	10.75.241.1	53420	54.215.12..	2200	TCP	182	41	8	5433	36	5469	209..	72	139776	0.009002000	0xd9f7	(53319)	4619	
4629	2022/336	19:16:06.396916	10.75.241.1	53420	54.215.12..	2200	TCP	78	41	8	5469	0	5469	209..		142336	0.002000000	0xd9f8	(53320)		
4630	2022/336	19:16:06.396916	10.75.241.1	53420	54.215.12..	2200	TCP	78	41	8	5469	0	5469	209..		144384	0.000000000	0xd9f9	(53321)		
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	0xd9e3	(53313)	4628	
4634	2022/336	19:16:07.509904	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	289..	1448	22480	5469	7240	66680	1.203024000	0xd120	(53542)		
4641	2022/336	19:16:09.799904	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	289..	1448	22480	5469	7240	66680	2.209044000	0xd227	(53799)		
4652	2022/336	19:16:21.000920	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	289..	1448	22480	5469	7240	66680	4.209012000	0xd22c	(53811)		
4659	2022/336	19:16:23.702221	10.75.241.1	53420	54.215.12..	2200	TCP	72	41	8	5469	0	5469	209..		144384	7.702154000	0xd2fa	(53322)		
4668	2022/336	19:16:23.722221	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	224..	1448	23856	5469	7240	66680	0.019000000	0xd6d6	(58098)		
4681	2022/336	19:16:22.224221	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	289..	1448	22480	5469	7240	66680	0.493018000	0xd7fc	(53852)		
4721	2022/336	19:16:36.818922	10.75.241.1	53420	54.215.12..	2200	TCP	78	41	8	5468	0	5468	209..		144384	14.596791000	0xd9f6	(53323)		
4732	2022/336	19:16:36.818922	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	224..	1448	23856	5469	7240	66680	0.020018000	0xd9d6	(53854)		
4738	2022/336	19:16:38.403554	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	289..	1448	22480	5469	7240	66680	1.567931000	0xd9df	(57823)		
6434	2022/336	19:16:53.938824	10.75.241.1	53420	54.215.12..	2200	TCP	78	41	8	5468	0	5468	209..		144384	13.589269000	0xd9fc	(53324)		
6435	2022/336	19:16:53.940824	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	224..	1448	23856	5469	7240	66680	0.036081000	0xd9f3	(57780)		
8824	2022/336	19:17:07.483123	10.75.241.1	53420	54.215.12..	2200	TCP	78	41	8	5468	0	5468	209..		144384	15.001701000	0xd9f6	(53325)		
8850	2022/336	19:17:07.833123	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	224..	1448	25304	5469	7240	66680	0.020000000	0xf22c	(62118)		
8886	2022/336	19:17:07.835123	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	253..	1448	26752	5469	7240	66680	0.000000000	0xf229	(62121)		
9436	2022/336	19:17:18.686196	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	389..	1448	22480	5469	7240	66680	3.572871000	0xf73f	(62452)		
12945	2022/336	19:17:22.118422	10.75.241.1	53420	54.215.12..	2200	TCP	78	41	8	5468	0	5468	209..		144384	11.512238000	0xd9f6	(53326)		
12946	2022/336	19:17:22.118422	10.75.241.1	53420	54.215.12..	2200	TCP	1514	64	8	224..	1448	23856	5469	7240	66680	0.016000000	0xf22c	(62120)		

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.

- Deactivate and then reactivate the outbound SSH, as shown below:

```
user@switch# deactivate system services outbound-ssh client mist
```

```
user@switch# activate system services outbound-ssh client mist
user@switch# commit
```

Watch the following video as well for more information on how to troubleshoot a switch:



Video: [Wired Assurance Troubleshooting](#)

## Troubleshooting Juniper CloudX

This chapter takes you through the steps involved in checking if your switch is communicating with Mist cloud by using CloudX.

To check if a switch communicates with Mist cloud using CloudX:

1. Run the below CLI commands on the switch:

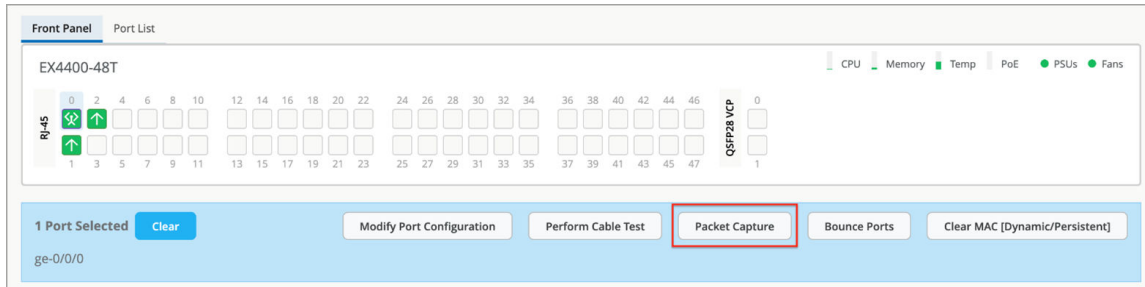
```
{master:0}
user@switch> show version | match mist
JUNOS Mist Agent [v1.0.2205-2]

{master:0}
user@switch> show system connections | grep 443
tcp4      0      0 192.168.2.52.62957
52.52.102.40.443          ESTABLISHED
```

To verify CloudX through the Mist portal, you can use the steps below:

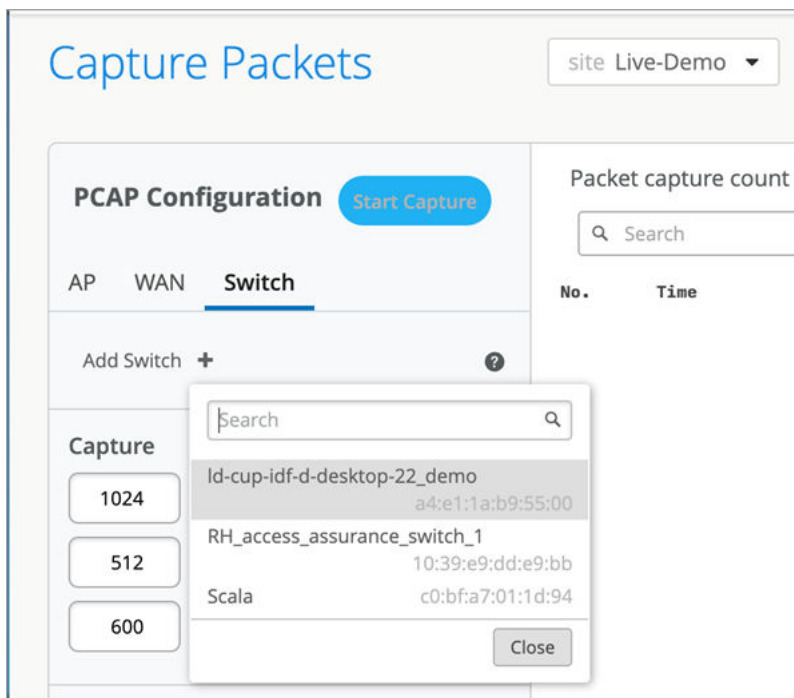
- a. Log in to the Mist portal ([manage.mist.com](https://manage.mist.com)).
- b. Click **Switches** > *switch name* to go to the switch detail page.
- c. Click any port or a range of ports.
 

If CloudX is running, the **Packet Capture** button is enabled; otherwise, the button is grayed out.



You can also check if CloudX is enabled on multiple switches by using the Mist portal.

To do that, click **Site > Switch Packet Captures > Add Switch**.



The switches listed here are all CloudX-enabled.

2. Verify that Mist Cloud Daemon (mcd) and Junos Mist Daemon (jmd) are running.

mcd is responsible for enabling communication between the switch and the cloud. It maintains a secure WebSocket connection to the terminator in the cloud.

jmd is used for:

- Generating periodic statistics for the device.
- Applying device configuration.
- Gathering device events.
- Initiating device functions (such as packet capture and software updates).

- Returning results from requested functions (such as files and streamed data).

To verify that jmd and mcd are running, use the following CLIs:

```

user@switch> start shell
% ps aux | grep jmd
root  21408  0.0  0.4 1246080 32200  - S   Fri23    15:17.51 /var/run/scripts/jet/jmd -
mcd-socket /var/run/mist_mcd.ipc
mist  3706  0.0  0.0  11136  2516  0 S+  07:14    0:00.00 grep jmd
%
%
% ps aux | grep mcd
root  21319  0.0  0.3 1242924 22256  - I   Fri23    8:18.00 /var/run/scripts/jet/mcd
root  21408  0.0  0.4 1246080 32200  - S   Fri23    15:17.53 /var/run/scripts/jet/jmd -
mcd-socket /var/run/mist_mcd.ipc
mist  3708  0.0  0.0  11136  2516  0 S+  07:14    0:00.00 grep mcd
%

```

3. Check the jmd and mcd logs for any errors by using the CLI commands below. Typically, jmd logs shows issues related to configuration or stats. The mcd logs report issues related to the connectivity between the switch and the cloud.

```

user@switch> show log jmd.log | last 10
[jmd] 2024/11/04 07:12:02 collector.go:850: total stats collection time = 10s
[jmd] 2024/11/04 07:12:02 app_states.go:355: app sending stats to mist cloud (26171 bytes)
[jmd] 2024/11/04 07:12:02 app_states.go:360: successfully sent ipc stats:
[jmd] 2024/11/04 07:12:02 app.go:282: processing app state "STEADY"
[jmd] 2024/11/04 07:12:12 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:22 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:32 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:42 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:52 app.go:339: sending ipc keep-alive
[jmd] 2024/11/04 07:12:52 collector.go:417: collecting periodic stats, interval 60

```

```

user@switch> show log mcd.log | last 10
[mcd] 2024/11/04 07:09:31 app.go:967: successfully sent msg to cloud: ep-telemetry
[mcd] 2024/11/04 07:10:02 ipc_server.go:414: rx ipc request: send cloud telemetry
[mcd] 2024/11/04 07:10:02 ipc_server.go:447: forwarding ipc telemetry to "junos-stats-"

```



```
(26167 bytes)
[mcd] 2024/11/04 07:11:02 ipc_server.go:414: rx ipc request: send cloud telemetry
[mcd] 2024/11/04 07:11:02 ipc_server.go:447: forwarding ipc telemetry to "junos-stats-"
(26171 bytes)
[mcd] 2024/11/04 07:12:01 app.go:967: successfully sent msg to cloud: ep-telemetry
[mcd] 2024/11/04 07:12:02 ipc_server.go:414: rx ipc request: send cloud telemetry
[mcd] 2024/11/04 07:12:02 ipc_server.go:447: forwarding ipc telemetry to "junos-stats-"
(26171 bytes)
[mcd] 2024/11/04 07:13:02 ipc_server.go:414: rx ipc request: send cloud telemetry
[mcd] 2024/11/04 07:13:02 ipc_server.go:447: forwarding ipc telemetry to "junos-stats-"
(26171 bytes)
```

4. If jmd or mcd is not running for some reason, try restarting it, as shown in the sample below.

```
{master:0}
user@switch> request extension-service restart-daemonize-app mcd
Extension-service application 'mcd' with pid: 92502 exited with return: -1
Extension-service application restarted successfully
```



**NOTE:**

```
request extension-service daemonize-restart mcd
```

5. If the switch is not connecting to the cloud, check its reachability by using a ping and curl test. These tests will help you check if the required firewall ports are allowed.

The cloud endpoints are not set up to respond to ping tests; however, running a ping test will ensure that DNS resolves FQDN. Here is a sample ping test:

```
user@switch> ping jma-terminator-staging.mistsys.net
PING a8481a00030ad459aac15af07d5f2c5b-75855524.us-east-1.elb.amazonaws.com (3.210.247.53): 56
data bytes
^C
--- a8481a00030ad459aac15af07d5f2c5b-75855524.us-east-1.elb.amazonaws.com ping statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
```

Here is a sample curl test:

```
mist@Scala> start shell
```

```
% curl -k https://jma-terminator-staging.mistsys.net/about
{
  "version": "0.3.14906",
  "git-commit": "2ec94c073f64a182fe927c4037d871a8d58c1149",
  "build-time": "2024-11-05T04:51:17Z",
  "go-runtime": "go1.23.2",
  "env": "staging",
  "procname": "jma-terminator/ /provider=aws/env=staging/host=ip-172-31-58-252/pid=1/
user=root",
  "start-time": "2024-11-05T05:05:29Z",
  "uptime": 37341.554596,
  "private-instance": false
}
```

A valid response from the curl test proves that the jma-terminator in the Mist cloud is reachable. A lack of response or receipt of an error will indicate that the path between the switch and the cloud is blocking these ports, likely because of the firewall. The URLs used in the test are the same as those in [firewall ports](#) and differ between cloud instances.

## RELATED DOCUMENTATION

| [Juniper CloudX Overview](#) | 21

# Cloud-Ready LED Blink Patterns

## IN THIS SECTION

- [Cloud-Ready Connection Process](#) | 274

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 19: Cloud LED Blink Patterns**

CLD LEDs	Blink Pattern	Meaning
•	solid green	The ZTP process is complete.
○	solid white	Connected to Mist cloud.
• • •	3 yellow	No IP Address. The DHCP server is not configured or could not be reached. Junos did not receive a DHCP lease or IP address.
• • • •	4 yellow	No default gateway. Either the address was not received or it is not configured on the device,
• • • • •	5 yellow	The default gateway could not be reached. No ARP from the default gateway.
• • • • • •	6 yellow	No DNS server(s) found in the static configuration, or in the DHCP lease.

Table 19: Cloud LED Blink Patterns (Continued)

CLD LEDs	Blink Pattern	Meaning
	7 yellow	No response from the DNS server. The switch received an IP address for the DNS server via DHCP, but it cannot not reach the Mist cloud.
	9 yellow	The Mist agent cannot reach the Mist cloud.
	1 yellow, pause, 2 yellow	Could not connect to the redirect server, most likely due to a firewall blocking TCP port 443, TCP port 2200. See also <a href="#">Ports to Open in Your Firewall</a> .
	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 <b>redirect.juniper.net</b> .
	1 yellow, pause, 5 yellow	Incorrect time on the switch. During ZTP, the phone home client (PHC) received a certificate with the wrong time. ZTP could not continue.

Table 19: Cloud LED Blink Patterns (*Continued*)

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

## Cloud-Ready Connection Process

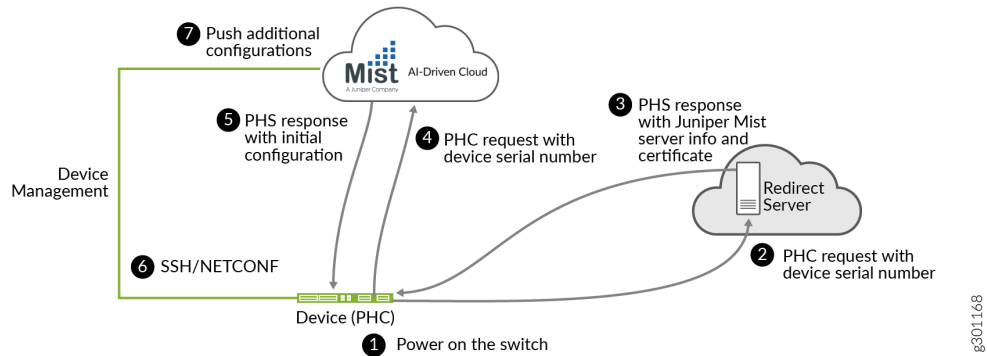
Juniper EX Series switches are cloud-ready devices, which means they are Day-0 capable of connecting to the Juniper Mist cloud. When running a supported version of Junos, these switches can also automatically establish a connection to Juniper Mist cloud services, where they can then be on-boarded (Day 1), managed (Day 2), and monitored (Day2+) from the Mist portal.

As part of zero-touch-provisioning (ZTP), a secure TCP connection uses pre-shared keys on both the device and cloud to establish a connection. Figure 1 provides a break-down of the ZTP process.

On the front of EX4100, EX4100-F, and EX4400 switches there is a CLD interface LED that you can use to monitor the ZTP progress. The blink pattern of the LED can help troubleshoot any Day 0 connection issues.

For switches already in the cloud, you can view the connection status from the Switches page of the Mist portal.

Figure 25: Stages in the ZTP Process for Cloud Ready Switches



The first time a cloud-ready switch is powered on, an on-board phone-home client (PHC) connects to a redirect server, which then redirects it to a phone home server (PHS) where the switch can get the latest Junos configuration. You can also have the switch connect to a DHCP server that supports ZTP and run the ZTP process from there.

#### RELATED DOCUMENTATION

| [Cloud-Ready LED Blink Patterns](#) | 271

#### RELATED DOCUMENTATION

| [Cloud-Ready Connection Process](#) | 274

## 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? | 276](#)
- [How to check which VLAN is missing on the switch port? | 277](#)
- [How to verify if Marvis is detecting the correct case of missing VLANs? | 277](#)
- [How to fix the missing VLAN error? | 278](#)
- [In an EX Series switch-based Virtual Chassis \(VC\) system, how do I convert the DAC-attached VC port to a trunk port? | 278](#)
- [Is Spanning Tree Protocol \(STP\) enabled on EX Series switches by default? | 278](#)
- [I created a port configuration and applied that to a range of ports on a switch. Now, I want to select a single port from that range and configure an individual port description on it. How do I do that? | 279](#)
- [Upgrade Required: Upgrading Junos or the BIOS for EX4400 Switches From the Mist Dashboard | 279](#)

## 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.

## 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.

**ACTIONS**

3 Switch

0 Authentication

0 DHCP, DNS

RF

Security

AP

Missing VLAN

Bad Cable

Negotiation Mon...

**MISSING VLAN**

Switch (3)

**RECOMMENDED ACTION**

MISSING VLAN  
The below switches have 1 or more VLANs missing on the ports where Mist APs are connected to. Please add the VLANs to the respective switch ports.

ID	Site	Switch	Details	Date
MV-16454	ARB1	ARB-SW-2-IDF-03	VLANs (260, 261, 262) missing	Feb 25, 2020 12:13 PM
MV-16455	ORD	ORD-SW-MDF1-DIS1	VLAN (260) missing	Mar 22, 2020 05:59 AM
MV-16456	ARB1	ARB-SW-IDF-2-02	VLANs (58, 260, 261, 262) missing	Feb 25, 2020 12:14 PM

▼ STATUS

**LATEST UPDATES**

Filter by action type

Today

3/28/2020, 12:34:14 PM  
AI VALIDATED  
Missing VLAN (ID: MV-16452)  
Site: ARB1  
Switch: ARB-SW-2-IDF-03  
Port: G11/0/43  
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: G11/0/4  
AP: ARB1-AP-LOBBY  
VLAN: 58, 261, 262

Yesterday

3/27/2020, 3:25:45 AM  
AI VALIDATED  
Missing VLAN (ID: MV-16172)  
Site: ARB1  
Switch: ARB-SW-IDF-2-02  
Port: G11/0/39  
AP: ARB1-AP-218  
VLAN: 58, 260

3/27/2020, 3:25:45 AM  
AI VALIDATED  
Missing VLAN (ID: MV-16173)  
Site: ARB1

## How to verify if Marvis is detecting the correct case of missing VLANs?

To verify whether the Marvis AI is detecting the correct case of missing VLANs, do a packet capture or port mirroring on the switch port to which the AP is connected, and use the Wireshark tool to analyze the traffic. You can also use the VLAN filter to verify if any traffic is coming from that VLAN. See [Dynamic and Manual Packet Captures](#) for more help on setting up Wireshark.



## How to fix the missing VLAN error?

Once you have identified the VLAN that is missing from the switch port but is being used in your AP or WLAN configuration, you can configure that VLAN on your switch. After the VLAN is correctly configured on your switch and the AP starts detecting traffic from it, Mist takes some time to verify the fix and ensure that the issue is resolved. After that, the warning disappears automatically.



**NOTE:** If you see the warning even after fixing all the VLANs on your switch ports, open a support ticket for assistance. For more information, see [Create a Support Ticket](#).

## In an EX Series switch-based Virtual Chassis (VC) system, how do I convert the DAC-attached VC port to a trunk port?

You can do this by converting the VC port on the Virtual Chassis to network port by using the following CLI commands in the Additional CLI Commands section on the switch details page (**Switches > Switch Name**):

```
request virtual-chassis mode network-port
```

After saving the above configuration, reboot the switch using the Reboot option on the Utilities menu on the switch details page (**Switches > Switch Name**). This behavior applies to the EX Series switches that come with a default VC port.

## Is Spanning Tree Protocol (STP) enabled on EX Series switches by default?

STP is not enabled on EX Series switches by default. However, if the switch is managed by Mist, this protocol is enabled on the switch automatically. If users do not want STP on a port, they can set the port as a Rapid Spanning Tree Protocol (RSTP) edge port by using a port profile. This setting ensures that the port is treated as an edge port and guards against the reception of Bridge Protocol Data Units (BPDUs).

**I created a port configuration and applied that to a range of ports on a switch. Now, I want to select a single port from that range and configure an individual port description on it. How do I do that?**

To do that, follow the steps below:

1. Go to the switch details page (**Switches** > *Switch Name*).
2. Select the port to be updated from the Front Panel section.
3. Click the Modify Port Configuration button to open the Port Configuration tile.
4. Specify a port description in the Description field.
5. Click the check mark (✓) on the top right of the Port Configuration tile to save the configuration.



**NOTE:** Saving this port configuration will remove the overlapping port IDs from the existing configurations (applied to the range of ports) and create a new port configuration for the selected port.

## **Upgrade Required: Upgrading Junos or the BIOS for EX4400 Switches From the Mist Dashboard**

For Juniper EX Series switches that are managed by Mist, you can identify and upgrade the Junos version and/or switch BIOS to the current version from the Mist dashboard. To alert you to a version mismatch and to recent releases, a notification appears at both the switch list, and on the configuration page of individual switch. When more than one switch is affected, you can select all to perform a simultaneous upgrade.

Figure 26: Upgrade Notification

The screenshot shows the Mist Switches management interface for site 'Bangalore-site-2'. The interface displays various metrics: 5 Adopted Switches, 0 Discovered Switches, 5 Wired Clients, and 0 W Total Allocated AP Power. Below these metrics are buttons for 'Switch-AP Affinity', 'PoE Compliance', 'VLANs', '83% Version Compliance', 'Switch Uptime', and '60% Config Success'. A table lists the switches with columns for Status, Name, IP Address, Model, Mist APs, Wireless Clients, Wired Clients, MAC Address, and Version. One switch, 'Dist2' (IP: 172.31.14.37, Model: EX4400-24P), has a 'BIOS Upgrade Required' notification. A link at the bottom suggests 'Assign Switches to the site from the Inventory page'.

Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients	MAC Address	Version
Connected	Access	172.31.14.38	EX4100-F-12P	0	0	--	...	22.4R1.10
Connected	CORE1	172.31.14.24	QFX5130-32CD	0	0	2	...	22.4R1.11
Connected	core2	172.31.14.33	QFX5130-32CD	0	0	--	...	22.4R1.11
Reboot to use new image	Dist1	172.31.14.36	EX4400-48MP	0	0	3	...	22.4R2.8
BIOS Upgrade Required	Dist2	172.31.14.37	EX4400-24P	0	0	--	...	22.4R2-S1

When you see an upgrade notification, we recommend that you perform the update as soon as convenient to prevent "silent," or unexpected, device reboots that can result in the case of new Mist features not being supported by a previous version of the switch firmware. Note that for BIOS upgrades and some Junos upgrades, the switch will be taken offline to reboot or restart a given line card.

The **Upgrade BIOS** option appears:

The screenshot shows the 'Upgrade BIOS' dialog box in the Mist Switches management interface. The dialog box contains the following information: 'Total Switches selected to upgrade: 1', 'Switch Model: EX4400-24P', 'Selected Switches: 341143700', and 'Upgrade to Version: CDEN\_P\_EX1\_00.20.01.00'. There are two checked options: 'Reboot switch after image copy' and 'I accept the End User License Agreement'. The dialog box has 'Start Upgrade' and 'Cancel' buttons. In the background, the 'Upgrade BIOS' button is highlighted in red in the top right corner of the interface.

For switches configured for virtual chassis (VC), all members are automatically upgraded together; you do not need to perform separate BIOS upgrades. Likewise, in the case where different FPCs are running different BIOS versions, only those FPCs that need the newer BIOS will be upgraded.

Tip: You can find the switch BIOS by opening a remote shell (**Utilities > Remote Shell** at the top of the switch configuration page) and running the following Junos command from the CLI:

```
show chassis firmware
```

The exact output will depend on the switch and hardware configuration, as shown in the following sample:

```
{master:0} user@device> show chassis firmware
          Part      Type      Version
FPC 0    loader     FreeBSD EFI loader 2.1
          BIOS      CDEN_P_EX1_00.15.01.00
          System CPLD 0.f
          CPU CPLD  1.0
FPC 1    loader     FreeBSD EFI loader 2.1
          BIOS      CDEN_P_EX1_00.20.01.00
          System CPLD 0.12
          CPU CPLD  1.0
{master:0} user@device>
```

The following tables provide upgrade recommendations with regards to a known, and since September, 2023 fixed, issue with silent reboots on the EX4400 platform. If you have a device that is effected, upgrade the BIOS and/or Junos version as instructed in the following tables.

**Table 20: Junos Upgrade Warning**

Junos Version	Next Step	Comment
Junos 22.2R3-S1.11 and below.	Upgrade to Junos 22.2R3-S2.	Schedule a maintenance window and perform the upgrade. Multiple EX4400s can be upgraded together. Go to <b>Utilities &gt; Upgrade Firmware</b> on the Switch detail page to perform the upgrade.
Junos 22.3	Upgrade to Junos 22.3R2-S2.	
Junos 22.4 and below 22.4R2-S1	Upgrade to Junos 22.4R2-S1.	

**Table 21: BIOS Upgrade Warning**

BIOS Upgrade Required	Next Step	Comment
-----------------------	-----------	---------

If a warning for “BIOS Upgrade Required” is displayed on the Mist UI...	<p>Upgrade the affected switches to the recommended BIOS version.</p> <p>If the switch is configured for VC, and is operating in <a href="#">HiGiG mode</a> network-port or HGoE mode you will need to manually re-enable that virtual-chassis mode on the device and reboot. No such action is needed if the switch is running the default VC mode, HiGiG mode.</p>	<p>Upgrading the BIOS upgrade requires the switch to reboot. Only those versions needing the upgrade will be upgraded, for example in a virtual chassis.</p> <p>On the Switch Detail page, go to <b>Utilities &gt; Upgrade BIOS</b>. Select the EX4400 switches with the warnings and click <b>Upgrade BIOS</b>.</p>
	<p>To see which mode VC is running in, open a remote shell to the switch and run the following commands: <code>show virtual-chassis mode</code> and <code>show virtual-chassis vc-port</code></p>	<p>For additional details, see <a href="#">TSB71527</a>.</p>

**Table 22: BIOS and Junos Upgrade Warning**

Both Warnings	Next Step	Comment
If a warning for BIOS as well as Junos Upgrade Required is displayed on the Mist UI...	We recommend that you perform both upgrades and then reboot once to minimize downtime.	On the Switch List page, select the affected switches and click <b>Upgrade Switches (without Reboot)</b> and <b>Upgrade BIOS (with Reboot)</b> to upgrade both Junos and BIOS in a single reboot.

## RELATED DOCUMENTATION

[https://www.juniper.net/documentation/us/en/software/mist/mist-wired/topics/topic-map/manage-template-settings-wired-assurance.html#id\\_dd4\\_2sm\\_4yb](https://www.juniper.net/documentation/us/en/software/mist/mist-wired/topics/topic-map/manage-template-settings-wired-assurance.html#id_dd4_2sm_4yb)

# 8

CHAPTER

## Appendix

---

[Deploy and Manage EX Series Switches at a Branch | 284](#)

---

# Deploy and Manage EX Series Switches at a Branch

You can use Juniper Networks® EX Series Switches at branch locations as traditional standalone switches or as a Virtual Chassis combining up to ten switches and managing them as a single device. For higher scale deployments at the branch, you can use topologies that contains distribution switches between access switches and WAN devices.

The following document takes you through the workflow involved in deploying and managing EX Series Switches at the branch using the Juniper Mist™ cloud.

[Distributed Branch EX Series—Juniper Validated Design \(JVD\)](#).