Engineering
Simplicity

Virtual Chassis User Guide for Switches

Published
2023-05-17

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

Virtual Chassis User Guide for Switches

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

Table of Contents

About This Guide

Use this guide to set up and configure an EX2300, EX3400, EX4300, EX4400, EX4600, EX4650, or QFX Series Virtual Chassis. A Virtual Chassis is composed of a supported combination of multiple switches that operate and are managed as a single switch or network entity. Refer also to the hardware documentation for the types of switches comprising the Virtual Chassis for more details on how to physically interconnect them.

1

CHAPTER

Virtual Chassis Overview

Virtual Chassis Overview for Switches

IN THIS SECTION

Many Juniper Networks EX Series and QFX Series switches support *Virtual Chassis*, a flexible and scalable technology with which you can connect individual switches together to form one unit, and configure and manage the unit as a single chassis. Virtual Chassis ports (VCPs) connect member switches together to form a Virtual Chassis, and are responsible for passing all data and control traffic between member switches.

Use the following links to find the right Virtual Chassis User Guide for different EX Series and QFX Series switches if this guide doesn't cover the switches you're interested in:

- Virtual Chassis User Guide for EX2200, EX3300, EX4200, EX4500 and EX4550 Switches covers configuring and maintaining EX2200, EX3300, EX4200, EX4500, and EX4550 Virtual Chassis.
- Virtual Chassis User Guide for EX8200 Switches describes configuring and maintaining EX8200 Virtual Chassis.
- For deployments with EX9200 switches, you should plan or move to MC-LAG or Junos Fusion Enterprise architectures rather than using a Virtual Chassis. We do not recommend using EX9200 switches in a Virtual Chassis. If needed to aid in migration away from EX9200 Virtual Chassis, see Virtual Chassis User Guide for EX9200 Switches.
- Virtual Chassis User Guide for Switches describes configuring and maintaining *all other* EX Series and QFX Series Virtual Chassis.

NOTE: Virtual Chassis Fabric (VCF) is an evolution of Virtual Chassis technology. VCF configurations share some elements of Virtual Chassis technology, including configuration statements and administrative commands, but use a spine-and-leaf topology with supported QFX Series switches as spine devices and supported EX Series and QFX Series switches as leaf devices. For details on configuring and maintaining a VCF, see the Virtual Chassis Fabric User Guide.

Benefits of Virtual Chassis on Switches

- Simplifies configuration and maintenance: Multiple devices can be managed as a single device with the same or similar capabilities as the standalone device.
- Increases fault tolerance and high availability (HA): A Virtual Chassis can remain active and network traffic can be redirected to other member switches when a single member switch fails.
- Flattens your network and reduces networking overhead by allowing network devices to synchronize to one resilient logical device rather than to multiple physical devices.
- Enables a simplified Layer 2 network topology that minimizes or eliminates the need for loop prevention protocols such as Spanning Tree Protocol (STP).
- Provides a flexible model for expanding your network: You can easily add Virtual Chassis member switches to increase the number of access ports on your network to support more servers, computers, phones, or other devices with minimal complications to the existing network topology and switch configuration.

Virtual Chassis Basics on Switches

Virtual Chassis is a feature on Juniper Networks routing or switching devices that provides network resiliency in the form of redundant Routing Engines and network expansion flexibility with minimal impact to a configuration.

Virtual Chassis technology on switches enables you to interconnect supported combinations of EX Series and QFX Series switches into one logical device that you can configure and manage as a single unit. Switches interconnected into a Virtual Chassis are called Virtual Chassis member switches, identified by a member ID within the Virtual Chassis.

Virtual Chassis member switches are interconnected and communicate with each other using Virtual Chassis ports (VCPs).

Connecting Member Switches with Virtual Chassis Ports

A switch is not recognized by the Virtual Chassis as a member switch until it is interconnected with the primary or interconnected with an existing member of the Virtual Chassis using VCPs. EX Series and QFX Series switches that can be in a Virtual Chassis might support one or more of the following VCP options:

- Network or uplink ports that support the option to be configured as VCPs. Most switches support this VCP option.
- Ports that are configured as VCPs in the default factory settings. These ports can also be converted into and used as network ports instead of as VCPs, and converted back into VCPs again if needed.
- Dedicated VCPs, which are ports that can only function as VCPs. Only a few switches have dedicated VCPs.

Available VCP options vary among the different switch models. See "Virtual Chassis Port Options" on page 41 for a summary of the ports that are supported as VCPs on different switches.

When a port is set as a VCP, it cannot be used for any other purpose. If you want to use the port for another purpose, you must delete the VCP setting using the `request virtual-chassis vc-port` command. You can run this command directly on the member whose uplink VCP setting you want to delete or through the primary member of the Virtual Chassis configuration.

> **CAUTION**: Deleting a VCP in a Virtual Chassis configuration can cause the Virtual Chassis configuration to split. For more information, see "Understanding Split and Merge in a Virtual Chassis" on page 71.

If redundant VCP links of the same speed are connected between the same two member switches of a Virtual Chassis, the ports automatically form a VCP Link Aggregation Group (LAG) or bundle that distributes the inter-member VCP traffic load among them. See "Understanding Virtual Chassis Port Link Aggregation" on page 68 for details.

Virtual Chassis Configuration

You configure and manage nearly all aspects of an EX Series or QFX Series Virtual Chassis through the primary switch of the Virtual Chassis. However, you can also configure Virtual Chassis parameters when a switch is a standalone switch not interconnected with other members yet, because any switch that supports being in a Virtual Chassis is by default a single-member Virtual Chassis with member ID 0. Upon connecting the switch with others in a Virtual Chassis, any Virtual Chassis configuration statements and uplink Virtual Chassis port (VCP) settings previously configured on the standalone switch remain part of its configuration.

You can set up an EX Series or QFX Series Virtual Chassis using a nonprovisioned or preprovisioned configuration. If you want to deterministically control the roles and member IDs assigned to the member switches when creating and managing a Virtual Chassis, use a preprovisioned configuration, which distinguishes member switches by associating their serial numbers with the member ID.

When adding new member switches to a preprovisioned Virtual Chassis, you might be able to simplify the procedure by using the autoprovisioning feature, which automatically converts the interconnecting links into VCPs when cabling the new switch into the Virtual Chassis under certain conditions and configuration settings. See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45 for details.

Configuring Interfaces for a Virtual Chassis

The member ID of an EX Series or QFX Series Virtual Chassis member switch functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID as the *slot* element of the interface name.

The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the `ethernet-switching` family. To include the FPC in the default VLAN, add the `ethernet-switching` family to the configurations for those interfaces.

Mixed and Non-mixed EX Series and QFX Series Virtual Chassis

A Virtual Chassis might consist of all the same type of switches or different types of switches in supported combinations.

Some combinations of switches in a Virtual Chassis comprise a *mixed Virtual Chassis*, which contains member switches that have operational differences requiring the Virtual Chassis to be configured with a mixed mode setting that enables all of the member switches to inter-operate successfully.

Some combinations of different types or models of switches can inter-operate in a Virtual Chassis without requiring the Virtual Chassis to be configured in mixed mode, such as different switches that can run the same Junos OS software image.

See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details on the different combinations of switches supported in a Virtual Chassis.

Virtual Chassis Member Switch Roles

Member switches in an EX Series or QFX Series Virtual Chassis operate in either a *primary* Routing Engine role, *backup* Routing Engine role, or *linecard* role. For some mixed Virtual Chassis, the member switches in the Routing Engine role are recommended (and in some cases required) to be particular types or models of switches. Any switch supported in a Virtual Chassis can operate in the linecard role.

A standalone switch that supports Virtual Chassis is by default a single-member Virtual Chassis that is assigned member ID `0` and operates in the primary Routing Engine role as the primary of itself. When connected and configured into a Virtual Chassis with other member switches, the switch will be assigned a unique member ID and might take on a different role.

A nonprovisioned Virtual Chassis uses a primary-role election algorithm to select the member switches that assume the primary and backup roles if an existing member switch in the Routing Engine role fails or when new member switches are added. In a preprovisioned Virtual Chassis, you assign the roles to each member switch when forming the Virtual Chassis and adding or replacing member switches.

See the following for details on EX Series and QFX Series Virtual Chassis member switch roles:

- "Understanding Virtual Chassis Components" on page 35
- "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62

Global Management of Member Switches in a Virtual Chassis

The interconnected member switches in a Virtual Chassis operate and can be configured as a single network entity.

The serial console port and dedicated out-of-band management port on individual switches have global virtual counterparts when the switches are interconnected in a Virtual Chassis configuration. You can connect to the primary switch by connecting a terminal directly to the console port of any member switch. A *virtual management Ethernet (VME)* interface allows you to remotely manage the Virtual Chassis configuration by connecting to the out-of-band management port of any member switch through a single IP address. You can perform remote configuration and administration of all member switches of the Virtual Chassis configuration using the Junos CLI through the VME interface. See "Understanding Global Management of a Virtual Chassis" on page 64 for details.

When setting up a Virtual Chassis on EX Series switches that support the EZSetup script, you can run EZSetup once to specify the identification parameters for the primary, and these parameters implicitly apply to all member switches of the Virtual Chassis.

On switches that support the J-Web user interface, you can view the Virtual Chassis as a single device in J-Web and apply various device management functions to all member switches of the Virtual Chassis.

High Availability Using Redundancy

Interconnecting EX Series or QFX Series switches into a Virtual Chassis increases your network's high availability. A Virtual Chassis is more fault tolerant than a standalone switch because it can remain active, forward traffic and provide sub-second convergence in the case of a device or link failure.

Standalone switches that support only a single Routing Engine can operate with a primary and a backup Routing Engine when configured into a Virtual Chassis, and therefore support some high availability features that would otherwise not be available on the switch, such as Graceful Routing Engine Switchover (GRES) for hitless failover.

You can also increase fault tolerance within a Virtual Chassis by configuring other supported high availability features. For example, configuring Link Aggregation Group (LAG) bundles that include member links on different switches in the same Virtual Chassis enables traffic traversing the LAG to be redirected from a Virtual Chassis member switch that fails to links on another active Virtual Chassis member switch.

Adaptability as an Access Switch or Distribution Switch

A Virtual Chassis configuration supports a variety of user environments because it can be composed of different types of switches. You can select different switch models to support various functions. For example, you might set up one Virtual Chassis access switch configuration composed of full Power over Ethernet (PoE) models to support users sitting in cubicles equipped with PCs and Voice over IP (VoIP) phones. You could set up another Virtual Chassis configuration with partial PoE models to support the company's internal servers, and another to support the company's external servers. You can alternatively use a Virtual Chassis in a topology as a distribution switch.

Virtual Chassis Provisioning From the Factory-Default State Using the Phone-Home Client

Phone-home provisioning on a Virtual Chassis is a form of zero-touch provisioning (ZTP). With phone-home provisioning, when a device in the factory-default state boots up, a phone-home client (PHC) process automatically starts running on the device. The PHC gets bootstrapping information over the network from a central network management data source called the phone-home server (PHS), and installs the intended software image and configuration on the device without requiring any user interaction at the remote site.

The PHC also supports phone-home provisioning on some EX Series Virtual Chassis. Check Feature Explorer and search for **phone-home** to see the Virtual Chassis platforms that support phone-home provisioning. Phone-home provisioning on a Virtual Chassis is an extension of standalone device phone-home support.

See Provision a Virtual Chassis Using the Phone-Home Client for details about how the PHC works to provision a Virtual Chassis.

The PHS is usually part of a network management system (NMS) that supports phone-home provisioning. The network administrator enters the information that defines how the devices or Virtual Chassis at remote sites should be set up. An organization might have more than one PHS for redundancy.

Requirements for Phone-Home Provisioning to Work for a Virtual Chassis

The PHC only supports Virtual Chassis that meet the following conditions:

- The Virtual Chassis member devices have dedicated or default-configured VCPs.

  "Virtual Chassis Port Options" on page 41 explains the different VCP types and what's supported on different platforms.

- The Virtual Chassis members are all the same type of device.

  In other words, it isn't a mixed-mode Virtual Chassis. See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54.

- All member devices have the factory-default configuration.

  The PHC process only runs when a device or Virtual Chassis is in the factory-default state.

- The Virtual Chassis member are interconnected in a ring topology using only dedicated or default-configured VCPs.

  How To Enable Phone-Home Provisioning on a Virtual Chassis shows a sample EX4300 Virtual Chassis that is wired this way.

- At least one Virtual Chassis member has an active connection to the network and can access an available PHS that supports provisioning a Virtual Chassis.

  The connection to the PHS can be through the Virtual Chassis VME interface or any network-facing port on any Virtual Chassis member. See "Understanding Global Management of a Virtual Chassis" on page 64 for more about how the VME interface works.

RELATED DOCUMENTATION

Understanding EX Series Virtual Chassis

IN THIS SECTION

This topic introduces EX Series Virtual Chassis. An EX Series Virtual Chassis is a supported combination of interconnected EX2200, EX2300, EX3300, EX3400, EX4100, EX4100-F, EX4200, EX4400-24X, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4300, EX4400, EX4600, EX4500, EX4550, or switches operating as one logical device and managed as a single chassis. Switches in a Virtual Chassis are called *member switches*.

For information on EX8200 Virtual Chassis, see Virtual Chassis User Guide for EX8200 Switches.

Virtual Chassis Support on EX Series Switches

In an EX Series Virtual Chassis, you can interconnect standalone switches in the following combinations into one logical device, and manage the logical device as a single chassis:

- EX2200 Virtual Chassis, composed of up to four EX2200 switches.
- EX2300 Virtual Chassis, composed of up to four EX2300 switches or up to four EX2300 multigigabit model switches (EX2300-24MP, EX2300-48MP). Starting in Junos OS Release 18.4R1, you can also combine EX2300 multigigabit model switches with other EX2300 switches in the same Virtual Chassis, which operates as a non-mixed Virtual Chassis.
  On EX2300 switches, the Virtual Chassis feature requires license. See Understanding Software Licenses for EX Series Switches.
- EX3300 Virtual Chassis, composed of up to ten EX3300 switches.
- EX3400 Virtual Chassis, composed of up to ten EX3400 switches.
- EX4100 Virtual Chassis, composed of up to ten EX4100/EX4100-F switches (any EX4100 models, including EX4100 multigigabit models and the EX4100-F model)
- EX4200 Virtual Chassis, composed of up to ten EX4200 switches.
- EX4300 Virtual Chassis, composed of up to ten EX4300 switches, including multigigabit models (EX4300-48MP). An EX4300 Virtual Chassis operates as a non-mixed Virtual Chassis if it is composed of only EX4300 multigigabit model switches, or composed of any combination of any other EX4300 switches excluding the multigigabit models. An EX4300 Virtual Chassis operates as a mixed EX4300 Virtual Chassis if it is composed of EX4300 multigigabit model (EX4300-48MP) switches mixed with any other EX4300 model switches.
- EX4400 Virtual Chassis, composed of up to ten EX4400 switches (any EX4400 models, including EX4400 multigigabit models)
- EX4400-24X Virtual Chassis, composed of up to ten EX4400 switches (any EX4400 models, including EX4400 multigigabit models using HGoE mode).
- EX4500 Virtual Chassis, composed of up to ten EX4500 switches
- EX4550 Virtual Chassis, composed of up to ten EX4550 switches
- EX4600 Virtual Chassis, composed of up to ten EX4600 switches.
- EX4650 Virtual Chassis, composed of up to four EX4650-48Y switches.
  Starting in Junos OS Release 19.3R1, an EX4650-48Y Virtual Chassis can have up to two member switches.
  Starting in Junos OS Release 20.1R1, an EX4650-48Y Virtual Chassis can have up to four member switches.
- Mixed EX4200 and EX4500 Virtual Chassis, composed of up to ten total EX4200 and EX4500 switches

- Mixed EX4200 and EX4550 Virtual Chassis, composed of up to ten total EX4200 and EX4550 switches
- Mixed EX4200, EX4500, and EX4550 Virtual Chassis, composed of up to ten total EX4200, EX4500, and EX4550 switches
- Mixed EX4300 and EX4600 Virtual Chassis, composed of up to ten total EX4300 (excluding multigigabit models) and EX4600 switches. The member switches in the primary and backup Routing Engine roles must be EX4600 switches.
- Mixed EX4500 and EX4550 Virtual Chassis, composed of up to ten total EX4500 and EX4550 switches

Table 1 on page 11 and Table 2 on page 13 list the initial Junos OS release for combinations of switches interconnected into an EX Series Virtual Chassis. "N/A" indicates the combination is *not supported*. Switches must be running the same version of Junos OS software to join a Virtual Chassis, although the actual images might be different on different types of switches when mixed hardware models are supported together in a Virtual Chassis.

Table 1: Minimum Junos OS Release For EX Series Switch Combinations in a Virtual Chassis

Switch	EX2300 Switch	EX3400 Switch	**EX4100/ EX4100-F Switch**	EX4300 Switch	EX4400 Switch	**EX4400-24X Switch**	EX4600 Switch	EX4650
EX2300	15.1X53-D50, 18.1R2 for MP models, 18.4R1 for MP models combined with other models	N/A	N/A	N/A	N/A		N/A	N/A
EX3400	N/A	15.1X53-D50	N/A	N/A	N/A		N/A	N/A
EX4100/ EX4100-F	N/A	N/A	22.2R1	N/A	N/A		N/A	N/A

Table 1: Minimum Junos OS Release For EX Series Switch Combinations in a Virtual Chassis *(Continued)*

Switch	EX2300 Switch	EX3400 Switch	EX4100/ EX4100-F Switch	EX4300 Switch	EX4400 Switch	EX4400-24X Switch	EX4600 Switch	EX4650
EX4300	N/A	N/A	N/A	13.2X50-D10, 18.2R1 for MP models combined with other models	N/A		13.2X51-D25, excludes EX4300 MP models	
EX4400	N/A	N/A	N/A	N/A	21.1R1, 21.2R1 for MP models combined with other models		N/A	
EX4400-24X	N/A	N/A	N/A	N/A	N/A	23.1R1	N/A	
EX4600	N/A	N/A	N/A	13.2X51-D25, excludes EX4300 MP models	N/A		N/A	13.2X51-D25
EX4650	N/A	N/A	N/A	N/A	N/A		N/A	N/A

Table 2: Minimum Junos OS Release for Legacy EX Series Switch Combinations in a Virtual Chassis

Switch	EX2200 Switch	EX3300 Switch	**EX4100/ EX4100-F Switch**	EX4200 Switch	EX4500 Switch	EX4550 Switch
EX2200	12.2R1	N/A	N/A	N/A	N/A	N/A
EX3300	N/A	11.3R1	N/A	N/A	N/A	N/A
EX4100/ EX4100-F	N/A	N/A	22.2R1	N/A	N/A	N/A
EX4200	N/A	N/A	N/A	9.0R1	11.1R1	12.2R1
EX4500	N/A	N/A	N/A	11.1R1	11.1R1	12.2R1
EX4550	N/A	N/A	N/A	12.2R1	12.2R1	12.2R1

Basic Configuration of EX Series Virtual Chassis

Some EX Series switches can only form a Virtual Chassis with the same type of switches, while some can connect with other types of switches into a mixed-mode Virtual Chassis. See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details on the different types of switches that can be mixed in a Virtual Chassis.

You set up an EX Series Virtual Chassis by configuring Virtual Chassis ports (VCPs) on the member switches, and interconnecting the switches using the VCPs. VCPs are responsible for passing all data and control traffic between member switches in the Virtual Chassis. EX Series switches have one or more of the following VCP options:

- Network or uplink ports that you can configure into VCPs.
- Default-configured VCPs, which are configured as VCPs in the default factory configuration, but you can alternatively convert them into network or uplink ports if desired and reconvert them back into VCPs if needed.
- Dedicated VCPs, which you can only use as VCPs.

See "Virtual Chassis Port Options" on page 41 for details on which ports on different EX Series switches can be VCPs.

You can increase the VCP bandwidth between any two member switches by connecting multiple VCP links between the switches. When multiple VCP links interconnect the same two member switches, the links automatically form a Link Aggregation Group (LAG) bundle if they have the same speeds. For example, if you have four 40-Gbps links configured as VCPs between two member switches, the four links form a LAG with 160 Gbps of bandwidth. Similarly, connecting two 10-Gbps links configured as VCPs between two member switches creates a LAG with 2 member links at 20 Gbps total bandwidth. However, 10-Gbps and 40-Gbps links configured as VCPs between two Virtual Chassis member switches cannot be member links of the same VCP LAG.

Within a single wiring closet, you can add a new member switch to a Virtual Chassis by cabling the member switch into the Virtual Chassis using supported VCP links. You can also easily expand a Virtual Chassis configuration beyond a single wiring closet or over a longer distance by connecting member switches together using uplink ports that are supported as VCPs.

You can set up a Virtual Chassis using either a nonprovisioned or a preprovisioned configuration. If you want to deterministically control the role and member ID assigned to each member switch, use a preprovisioned configuration. Virtual Chassis member switches can have one of three roles: primary Routing Engine, backup Routing Engine, or linecard role. In some combinations of switches in a Virtual Chassis, we recommend or require that you configure certain switches into the Routing Engine roles. See "Understanding Virtual Chassis Components" on page 35 for more information about Virtual Chassis roles.

You can simplify adding switches to a preprovisioned configuration by using the automatic VCP conversion feature (see "Automatic Virtual Chassis Port (VCP) Conversion" on page 45), which automatically converts uplink ports into VCPs on the member switches on both sides of the new VCP links as you cable them. This Virtual Chassis expansion method is also called *autoprovisioning*.

EX2200 Switches in a Virtual Chassis

You can connect up to four EX2200 and EX2200-C switches into an EX2200 Virtual Chassis. EX2200 switches can’t be mixed in a Virtual Chassis with any other type of switches.

You can configure and use any EX2200 1-Gigabit Ethernet optical interfaces as VCPs. All RJ-45 interfaces, including built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors and 1000BASE-T RJ-45 transceivers, can also be configured into VCPs. You can configure up to eight 1-Gpbs interfaces configured as VCPs into a single Link Aggregation Group (LAG) bundle, so in an EX2200 Virtual Chassis, you can have a VCP LAG with a bandwidth of up to 8 Gigabit Ethernet.

EX2200 Virtual Chassis do not support other high availability features such as Graceful Routing Engine switchover (GRES), Nonstop bridging (NSB), Nonstop active routing (NSR), fast failover, and Nonstop software upgrade (NSSU).

You configure, monitor, and maintain an EX2200 Virtual Chassis is configured and managed similarly to other EX Series Virtual Chassis. See the following for details on configuring a Virtual Chassis with EX2200 switches:

- *Configuring an EX2200 Virtual Chassis (CLI Procedure)*

EX2300 Switches in a Virtual Chassis

You can connect up to four EX2300 switches into a Virtual Chassis. You can't combine EX2300 switches into a mixed Virtual Chassis with any other EX Series or QFX Series switches, but you can create a non-mixed Virtual Chassis with different models of EX2300 switches as follows:

- Any combination of up to four EX2300 and EX2300-C switches.
- Any combination of up to four EX2300 multigigabit model (EX2300-24MP and EX2300-48MP) switches.
- (Starting in Junos OS Release 18.4R1) Any combination of up to four EX2300 multigigabit model switches and other EX2300 or EX2300-C model switches. You do not need to set mixed mode.

NOTE: In Junos OS releases prior to 18.4R1, you cannot mix EX2300 multigigabit model switches with other EX2300 or EX2300-C model switches in an EX2300 Virtual Chassis.

EX2300 switches do not have default or dedicated VCP ports, but you can configure the 10-Gigabit Ethernet uplink ports as VCPs, and use those to interconnect the switches into a Virtual Chassis.

You connect EX2300 switches or EX2300 multigigabit switches into an EX2300 Virtual Chassis by configuring uplink ports as VCPs and using SFP+ transceivers. The uplink ports on EX2300 switches also support SFP transceivers, but you can't use SFP transceivers on uplink ports to form an EX2300 Virtual Chassis.

In all EX2300 Virtual Chassis, you can configure any EX2300 switch model into any member switch role (primary Routing Engine, backup Routing Engine, or linecard).

On EX2300 switches, the Virtual Chassis feature requires license. See Understanding Software Licenses for EX Series Switches.

You configure, monitor, and maintain an EX2300 Virtual Chassis in a similar way as other EX Series and QFX Series Virtual Chassis. See the following for details on configuring or changing the members in an EX2300 Virtual Chassis:

- "Configuring an EX2300, EX3400, or EX4300 Virtual Chassis" on page 80
- "Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis" on page 108
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

EX3300 Switches in a Virtual Chassis

You can interconnect up to ten EX3300 switches into a Virtual Chassis. EX3300 switches can't be mixed in a Virtual Chassis with any other type of switches.

Uplink ports 2 and 3 on EX3300 switches are configured as VCPs by default. You can change this default configuration or configure any of the other uplink ports as VCPs.

You configure, monitor, and maintain an EX3300 Virtual Chassis in a similar way as other EX Series Virtual Chassis. See the following for details on configuring or expanding a Virtual Chassis with EX3300 switches:

- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Adding a New Switch to an Existing EX3300 Virtual Chassis (CLI Procedure)*
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

EX3400 Switches in a Virtual Chassis

You can connect up to ten of any models of EX3400 switches into a Virtual Chassis. EX3400 switches can't be in a Virtual Chassis with any other EX Series or QFX Series switches.

By default, the QFSP+ uplink ports on EX3400 switches are configured as VCPs, or you can configure any SFP+ uplink module ports on EX3400 switches as VCPs. You can't use uplink ports with SFP tranceivers as VCPs to connect EX3400 switches into a Virtual Chassis.

You configure, monitor, and maintain an EX3400 Virtual Chassis in a similar way as other EX Series and QFX Series Virtual Chassis. See the following for details on configuring or changing the members in an EX3400 Virtual Chassis:

- "Configuring an EX2300, EX3400, or EX4300 Virtual Chassis" on page 80

- "Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis" on page 108
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

EX4100/EX4100-F Switches in a Virtual Chassis

You can interconnect up to ten EX4100 switches to form an EX4100 Virtual Chassis. All the member switches must be EX4100 switches (including EX4100-F and any other EX4100 multigigabit models); you can't mix EX4100 switches in a Virtual Chassis with any other EX Series switches.

You can use any of the EX4100 switch models in either the Routing Engine or linecard role in an EX4100/EX4100-F Virtual Chassis.

We have dedicated VCP ports available on EX4100/EX4100/F. You can't use any other ports on EX4100 switches as VCPs. The default VCPs are in PIC slot 1 on EX4100/EX4100-F switches, so the logical VCP interfaces are named vcp-255/1/0 through vcp-255/1/3.

If you don't need to use these ports as VCPs, you can convert them into network ports using the `request virtual-chassis mode network-port` command. The ports available in EX4100 can be converted into 25G/10G/1G network ports and the ports avaiable in EX4100-F can be converted into all 4x10G/4x1G network ports. However, no mixed mode is supported. To disable network port mode and return these ports to their default settings as VCPs, use the `request virtual-chassis mode network-port disable <reboot>` command. You must reboot the switch for a network port mode change to take effect.

You must use the `request virtual-chassis mode network-port` command to use the port as a network port. To disable the VCP port, use the command `request virtual-chassis vc-port set interface vcp-255/1/0 disable`

NOTE: You can't use one port as VCP and the other as a network port. The four ports must be both simultaneously used in VCP mode or network port mode.

Aside from the VCP differences from other switches, you configure, monitor, and maintain an EX4100/EX4100-F Virtual Chassis in a similar way as other EX Series and QFX Series Virtual Chassis. See the following for details on configuring and changing the members in an EX4100/EX4100-F Virtual Chassis:

- "Configuring an EX2300, EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis" on page 80
- "Adding a New Switch to an Existing EX2300, EX3400, EX4300, or EX4400 Virtual Chassis" on page 108
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

EX4200, EX4500, and EX4550 Switches in a Virtual Chassis

You can interconnect up to ten EX4200, EX4500, and EX4550 switches into a Virtual Chassis. You must set the Virtual Chassis to mixed mode if it is composed of EX4200 switches with EX4500 or EX4550 switches, but a Virtual Chassis with only one type of these switches or with only EX4500 and Ex4550 switches operates as a non-mixed Virtual Chassis.

EX4200 switches have two built-in dedicated VCPs. EX4500 or EX4550 switches have dedicated VCPs on the Virtual Chassis module, and the EX4500 or EX4550 switch must have the PIC mode set to Virtual Chassis mode to interconnect them into a Virtual Chassis. You can also configure any SFP, SFP+, and XFP uplink ports on all of these switches into VCPs. You can use the dedicated VCPs when the switches are close together, such as in the same wiring closet. Use uplink ports configured as VCPs for switches that located farther away, such as in different wiring closets.

You configure, monitor, and maintain an EX4200, EX4500, or EX4550 Virtual Chassis in a similar way as other EX Series Virtual Chassis. See the following for details on configuring a Virtual Chassis with these switches:

- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)*

For procedures on adding a new switch to a Virtual Chassis composed of any of these switches in a wiring closet, see:

- *Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis (CLI Procedure)*
- *Adding an EX4200 Switch to a Preprovisioned EX4500 Virtual Chassis or a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis (CLI Procedure)*
- *Adding an EX4500 Switch to a Preprovisioned EX4200 Virtual Chassis (CLI Procedure)*
- *Adding an EX4500 Switch to a Nonprovisioned EX4200 Virtual Chassis (CLI Procedure)*

EX4300 Switches in a Virtual Chassis

You can interconnect up to ten EX4300 switches to form an EX4300 Virtual Chassis. EX4300 switches can form a Virtual Chassis composed entirely of EX4300 switches (not multigigabit models) *or* entirely of EX4300 multigigabit model (EX4300-48MP) switches as a non-mixed Virtual Chassis.

You can also connect EX4300 multigigabit model switches and other EX4300 model switches together into an EX4300 Virtual Chassis by configuring all of the member switches into mixed mode. In this case, you must also configure the non-multigigabit EX4300 member switches with a special option (`ieee-`

`clause-82`) when setting mixed mode. The member switches in the Routing Engine role must be multitigabit model switches, and you must configure the non-multigigabit EX4300 switches into linecard role. Multigigabit model EX4300 switches can't be in a mixed Virtual Chassis with any other types of switches.

NOTE: If you remove an EX4300 member switch from a mixed EX4300 Virtual Chassis with multigigabit model members, you must disable `ieee-clause-82` port mode on the switch if you want to reconfigure it as a standalone switch or use it in any other type of mixed Virtual Chassis or any non-mixed Virtual Chassis. Otherwise, the VCPs will not connect with other members in the new Virtual Chassis.

You can have any EX4300 switches except multigigabit models in the linecard role in a mixed Virtual Chassis with EX4600 switches or QFX5100 switches. For more details on EX4300 member switches in a mixed Virtual Chassis, see "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54.

NOTE: EX4300 switches (excluding multigigabit models) can also be included as leaf nodes in a Virtual Chassis Fabric (VCF). For information on EX4300 switches in a VCF, see *Virtual Chassis Fabric Overview*.

On EX4300 switches excluding the multitigabit models, all 40-Gigabit Ethernet QSFP+ optical ports are configured as VCPs by default, and you can also configure any 10-Gigabit Ethernet uplink module ports into VCPs as needed. As a result, an EX4300 Virtual Chassis can have either 40-Gbps or 10-Gbps VCP links, or a combination of both. You can easily add new switches to a Virtual Chassis whether the switch is installed in the same building or at a different site because the ports available to be used as VCPs are long-distance optical ports.

On EX4300 multigigabit models, the four 40-Gigabit Ethernet QSFP+ ports on the rear panel are dedicated VCPs. These are the only ports on EX4300 multigigabit model switches that you can use as VCPs, so any EX4300 Virtual Chassis that contains EX4300 multigigabit model switches can have only have VCP links that are 40-Gbps.

You configure, monitor, and maintain an EX4300 Virtual Chassis in a similar way as other EX Series and QFX Series Virtual Chassis. See the following for details on configuring and changing the members in an EX4300 Virtual Chassis:

- "Configuring an EX2300, EX3400, or EX4300 Virtual Chassis" on page 80
- "Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis" on page 108

- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

EX4400 Switches in a Virtual Chassis

You can interconnect up to ten EX4400 switches to form an EX4400 Virtual Chassis. All the member switches must be EX4400 switches (including EX4400 multigigabit models or any other EX4400 models); you can't mix EX4400 switches in a Virtual Chassis with any other type of switches.

NOTE: Starting from Junos OS Release 23.1R1, EX4400-24X and EX4400 Switches can be connected as a single VC using HgoE mode.

By default, EX4400 switches come up in the default/HiGig mode and the Virtual Chassis ports come up as HiGig ports. To change this to HGoE mode, enter the `request virtual-chassis mode hgoe` command and reboot the switch, using the `request system reboot` command. To change this back to default/HiGig mode, enter the `request virtual-chassis mode hgoe disable` command, and reboot the switch using the `request system reboot` command.

NOTE: Whenever you change the mode of any existing stack from the default HiGiG to HgOE mode, note that the Virtual Chassis formation usually takes additional time to come up in HGoE mode. This is because it creates a fresh stack that results in additional convergence time.

You can use any of the EX4400 switch models in either the Routing Engine or linecard role in an EX4400 Virtual Chassis. You can use the default/HiGig mode or HGoE mode to achieve this:

NOTE: All the switches in a stack can operate in either the default/HiGig mode or HGoE mode. Mixing these two modes in a stack is not supported.

- **Virtual Chassis in default/HiGig mode:**

  On EX4400 switches, each of the two 100-Gigabit Ethernet QSFP28 ports on the rear panel are set as two logical 50G VCPs, forming four logical 50G VCP interfaces by default. The 100GbE ports can also accept 40GbE optics for Virtual Chassis connection or uplink connectivity. The default VCPs are in PIC slot 1 of EX4400 switches, so the VCP interfaces are named vcp-255/1/0, vcp-255/1/1, vcp-255/1/2, and vcp-255/1/3 respectively.

  To use 2X100G ports as network ports, you need to change the mode to Network port mode. When you switch to Network port mode, Virtual Chassis does not get created and rear panel ports work as 2X100G network ports. To change to Network port mode, enter the `request virtual-chassis mode network-port` command and reboot the switch, using the `request system reboot` command. To change back to default/HiGig mode, enter the `request virtual-chassis mode network-port disable` command and reboot the switch, using `request system reboot` command. Now, it gets switched to default/HiGig mode.

NOTE: In the default/HiGig mode, all ports of PIC slot 1 come up as Virtual Chassis ports. In the Network port mode, all ports of PIC slot 1 come up as Network ports.

If you disable one of these ports as a VCP using the `request virtual-chassis vc-port set interface vcp-255/1/0 member 1 disable command` command, that action disables the port as a VCP. Also, disabling vcp-255/1/0 disables both logical ports 0 and 1 (vcp-255/1/0 and vcp-255/1/1), and disabling vcp-255/1/2 disables both logical ports 2 and 3 (vcp-255/1/2 and vcp-255/1/3).

NOTE: You can't use one 100G/40G Ethernet port as a Virtual Chassis port and the other 100G/40G Ethernet port as Network port. The two ports must be both used simultaneously in Virtual Chassis port mode or Network port mode. When you insert a 40G transceiver, one `ifd` per physical port is enabled and the other `ifd` gets disabled. For the physical port 0, `vcp-255/1/0` is enabled and `vcp-255/1/1`is disabled. Similarly, for the physical port 1, `vcp-255/1/2` is enabled and `vcp-255/1/3`gets disabled."

- **Virtual Chassis in HGoE mode:**

  Starting with Junos OS Release 22.3R1, the HiGig protocol packets are encapsulated in a standard ethernet frame before these are sent out on the Virtual Chassis interface. This feature is known as HiGig over Ethernet (HGoE). To enable HGoE mode, use the `request virtual-chassis hgoe <reboot>` command. In HGoE mode, the 2X100G/40G ports available in the rear panel come up as Virtual Chassis ports by default.

  The 2x100G/40G interfaces or the optional 4x25G extension module interfaces can be configured as Virtual Chassis interfaces or network interfaces without rebooting the switch again. Use the `request virtual-chassis vc-port set pic-slot 1 port <port>` command to convert the interface to a Virtual Chassis port and the `request virtual-chassis vc-port delete pic-slot 1 port <port>` command to convert the interface to a network port . You can also also convert 4x25G network ports of the uplink module in PIC-slot 2 as Virtual Chassis ports. To convert the 4X25G network ports to Virtual ports and vice versa, use the `request virtual-chassis vc-port set pic-slot 2 port <port>` and `request virtual-chassis vc-port delete pic-slot 2 port <port>` commands.

  NOTE: In HGoE mode, on rear-panel ports (PIC slot 1) only 100G and 40G are supported as Virtual Chassis ports. On the the front panel, 4x25G uplink module ports (PIC slot2) 25G ports are supported as Virtual Chassis Ports. 10G speed/10-Gbps uplink module ports are not supported as Virtual Chassis ports.

Aside from the VCP differences from other switches, you configure, monitor, and maintain an EX4400 Virtual Chassis in a similar way as other EX Series and QFX Series Virtual Chassis. See the following for details on configuring and changing the members in an EX4400 Virtual Chassis:

- "Configuring an EX2300, EX3400, EX4300, or EX4400 Virtual Chassis" on page 80
- "Adding a New Switch to an Existing EX2300, EX3400, EX4300, or EX4400 Virtual Chassis" on page 108
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

EX4400-24X Switches in a Virtual Chassis

Starting with Junos OS Release 23.1R1, EX4400-24X Switches support Virtual Chassis. Note that the support is only on the HGoE mode and these switches do not support HiGiG mode. EX4400-24X swiches support Virtual Chassis functionality on 2x100G front panel ports, 4x25G uplink module, and the 1x100G uplink module. By default, none of the ports operate in the Virtual Chassis mode mode when the system comes up initially, and the 2x100G front panel ports operate as network ports by default.

NOTE: Starting from Junos OS Release 23.1R1, EX4400-24X and EX4400 Switches can be connected as a single VC using HgoE mode.

To enable the Virtual Chassis ports to get created on the 2x100G front panel ports, use the `request virtual-chassis hgoe <reboot>` command. The 2x100G interfaces or the optional 4x25G extension module interfaces can be configured as Virtual Chassis interfaces or network interfaces without rebooting the switch again. Use the `request virtual-chassis vc-port set pic-slot 1 port <port>` command to convert the interface to a Virtual Chassis port and the `request virtual-chassis vc-port delete pic-slot 1 port <port>` command to convert the interface to a network port . You can also also convert 4x25G network ports of the uplink module in PIC-slot 2 as Virtual Chassis ports. To convert the 4X25G network ports to Virtual ports and vice versa, use the `request virtual-chassis vc-port set pic-slot 2 port <port>` and `request virtual-chassis vc-port delete pic-slot 2 port <port>` commands.

NOTE: On the the front panel, 4x25G uplink module ports (PIC slot2) 25G ports are supported as Virtual Chassis Ports or network ports. 10G speed/10-Gbps uplink module ports are not supported as Virtual Chassis ports.

EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, and EX4400-48T Switches in a Virtual Chassis

Starting with Junos OS Release 23.1R1, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, and EX4400-48T Switches support Virtual Chassis. Note that the support is only on the HGoE mode and these switches do not support HiGiG mode. These swiches support Virtual Chassis functionality on 1x100G uplink module. By default, none of the ports operate in the Virtual Chassis mode mode when the system comes up initially, and the 1x100G uplink module port operates as a network port by default.

To enable the Virtual Chassis ports to get created on the 1x100G uplink module port, use the `request virtual-chassis hgoe <reboot>` command. The 1x100G uplink module port can be configured as Virtual Chassis interfaces or network interfaces without rebooting the switch again. Use the `request virtual-chassis vc-port set pic-slot 1 port <port>` command to convert the interface to a Virtual Chassis port and the `request virtual-chassis vc-port delete pic-slot 1 port <port>` command to convert the interface to a network port .

NOTE: 100G Network ports can be channelized into 4x25 Network ports. Channelization is not supported on Virtual Chassis ports.

EX4600 Switches in a Virtual Chassis

EX4600 switches can act as member switches in a non-mixed Virtual Chassis—a Virtual Chassis composed entirely of EX4600 switches—and also operate in a mixed Virtual Chassis with EX4300 switches. You can interconnect EX4300 switches (except multigigabit models, EX4300-48MP) with EX4600 switches as a mixed EX4600 Virtual Chassis.

You can interconnect up to ten EX4600 switches or a combination of up to ten EX4600 and EX4300 switches into an EX4600 Virtual Chassis. However, in a mixed Virtual Chassis with EX4600 and EX4300 switches, EX4600 switches must be in the primary and backup Routing Engine roles, and EX4300 member switches must be in the linecard role. As a result, at least two of the member switches in a mixed EX4600 Virtual Chassis must be EX4600 switches in the primary and backup Routing Engine roles. See "Understanding Virtual Chassis Components" on page 35 for more information about Virtual Chassis member roles.

EX4600 switches do not have any ports that are configured into VCPs by default, but you can configure any 40-Gigabit Ethernet QSFP+ and 10-Gigabit Ethernet SFP+ optical ports on an EX4600 switch into VCPs. Because EX4300 switches support similar VCP options, both non-mixed and mixed EX4600 Virtual Chassis can have 40-Gbps VCP links, 10-Gbps VCP links, or a combination of both.

You configure, monitor, and maintain an EX4600 Virtual Chassis in a similar way as other EX and QFX Series Virtual Chassis. See the following for more details on configuring and changing the members in an EX4600 Virtual Chassis, including a mixed EX4600 Virtual Chassis with EX4300 switches:

- "Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis" on page 93
- "Adding an EX4600 Switch to a Mixed or Non-mixed Virtual Chassis" on page 113
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

EX4650 Switches in a Virtual Chassis

Starting in Junos OS Release 19.3R1, you can interconnect up to two EX4650 switches in an EX4650 Virtual Chassis. The two member switches must be in the primary and backup Routing Engine roles.

Starting in Junos OS Release 20.1R1, you can interconnect up to four EX4650 switches in an EX4650 Virtual Chassis. You should configure two member switches into the primary and backup Routing Engine roles, and the remaining member switches into the linecard role.

See "Understanding Virtual Chassis Components" on page 35 for more information about Virtual Chassis member roles.

EX4650 switches can’t be combined with any other type of switches in a Virtual Chassis.

EX4650 switches do not have dedicated or default-configured VCPs, but you can set any of the 40-Gigabit Ethernet QSFP+ or 100-Gigabit Ethernet QSFP28 uplink ports on the front panel (non-channelized ports 48 through 55) as VCPs. You can’t use any of the other ports (network ports 0 through 47) as VCPs. Running the `request virtual-chassis vc-port set` command on the network ports doesn’t fail, but they will not function properly as VCPs.

An EX4650 Virtual Chassis operates the same as a QFX5120 Virtual Chassis, and you configure, monitor, and maintain it the same way as a QFX Series Virtual Chassis. See the following for more details on configuring and changing the members in an EX4650 Virtual Chassis:

- "Configuring an EX4650 or a QFX Series Virtual Chassis" on page 100

- "Adding a New Switch to an Existing EX4650 or QFX Series Virtual Chassis" on page 116
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

Release History Table

Release	Description
23.1R1	Starting in Junos OS Release 23.1R1, EX4400-24X Switches support Virtual Chassis.
22.3R1	Starting in Junos OS Release 22.3R1, the default Virtual Chassis ports for EX4400 can be converted to Network Ports and vice versa without the need for further reboot using the High Gigabyte Stacking over Ethernet (HGoE) mode.
20.1R1	Starting in Junos OS Release 20.1R1, an EX4650-48Y Virtual Chassis can have up to four member switches.
19.3R1	Starting in Junos OS Release 19.3R1, an EX4650-48Y Virtual Chassis can have up to two member switches.
18.4R1	Starting in Junos OS Release 18.4R1, you can also combine EX2300 multigigabit model switches with other EX2300 switches in the same Virtual Chassis, which operates as a non-mixed Virtual Chassis.

RELATED DOCUMENTATION

Virtual Chassis Overview for Switches | **2**

Understanding Virtual Chassis Components | **35**

Understanding Virtual Chassis Port Link Aggregation | **68**

Understanding QFX Series Virtual Chassis

IN THIS SECTION

This topic introduces QFX Series Virtual Chassis. A QFX Series Virtual Chassis is a supported combination of interconnected QFX3500, QFX3600, QFX5100, QFX5110, QFX5120 or EX4650, QFX5200, and EX4300 switches operating as one logical device and managed as a single chassis. Switches in a Virtual Chassis are called *member switches* of the Virtual Chassis.

This topic does not discuss Virtual Chassis Fabric (VCF). For information on understanding VCF, see *Virtual Chassis Fabric Overview*.

Virtual Chassis Support on QFX Series Switches

A QFX Series Virtual Chassis is a flexible, scaling switch solution for supported combinations of QFX3500, QFX3600, QFX5100, QFX5110, QFX5120, and QFX5200 switches. EX4650-48Y switches operate the same as QFX5120-48Y switches in a Virtual Chassis, so QFX Series Virtual Chassis configuration, monitoring and maintenance steps also apply to EX4650 Virtual Chassis. EX4300 switches can be included in some configurations of a QFX Series Virtual Chassis.

In a QFX Series Virtual Chassis, you can interconnect standalone switches in the following combinations into one logical device, and manage the logical device as a single chassis:

- Up to three QFX5200 switches (a non-mixed Virtual Chassis)
- Two QFX5120 switches or up to four EX4650 switches (a non-mixed Virtual Chassis), as follows:
  - Starting in Junos OS Release 19.3R1, you can interconnect two QFX5120-48Y or EX4650-48Y switches into a Virtual Chassis.
  - Starting in Junos OS Release 20.1R1, you can interconnect up to four EX4650-48Y switches into a Virtual Chassis.
  - Starting in Junos OS Release 20.2R1, you can interconnect two QFX5120-48T switches into a Virtual Chassis.

- Starting in Junos OS Release 20.3R1, you can interconnect two QFX5120-32C switches into a Virtual Chassis.
- Starting in Junos OS Release 23.1R1, you can interconnect four QFX5120-48YM switches into a Virtual Chassis.

- Up to ten QFX5110 switches or a combination of QFX5110 and QFX5100 switches (a non-mixed Virtual Chassis)
- Up to ten QFX5100 switches (a non-mixed Virtual Chassis)
- Up to ten QFX5100 switches with any combination of QFX3500, QFX3600, or EX4300 switches (a mixed-mode Virtual Chassis)
- Up to ten QFX3500 switches or QFX3600 switches, or a combination of QFX3500 and QFX3600 switches (a non-mixed Virtual Chassis)
- Up to ten QFX3500 or QFX3600 switches with supported EX4300 switches, or a combination of all three types of switches (a mixed-mode Virtual Chassis)

> **NOTE**: EX4300 multigigabit model (EX4300-48MP) switches are not supported in a mixed-mode Virtual Chassis with any QFX Series switches. Also note that the JNP-QSFP-100G-BXSR transceiver is not supported on Virtual Chassis Ports (VCP) connections.

Basic Configuration of QFX Series Virtual Chassis

Some QFX Series switches can only form a Virtual Chassis with the same type of switches, while others can be interconnected with other types of switches into a mixed-mode Virtual Chassis. See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details on the different types of switches that can be mixed in a Virtual Chassis.

You set up a Virtual Chassis by configuring Virtual Chassis ports (VCPs) on the member switches, and interconnecting the switches using the VCPs . VCPs are responsible for passing all data and control traffic between member switches in the Virtual Chassis.

The following ports on QFX Series switches that support Virtual Chassis can be configured into VCPs to form a QFX Series Virtual Chassis:

- On QFX5200 switches: Any 40-Gigabit Ethernet QSFP+ ports

NOTE: Starting in Junos OS Release 17.3R2-S4, 100-Gigabit Ethernet QSFP28 ports are also supported as VCPs on QFX5200 switches.

- On QFX5120 or EX4650 switches:
  - QFX5120-48Y or EX4650-48Y: Only the eight 40-Gigabit Ethernet QSFP+ or 100-Gigabit Ethernet QSFP28 front panel uplink ports (ports 48 through 55)
  - QFX5120-48T: Only the six 40-Gigabit Ethernet QSFP+ or 100-Gigabit Ethernet QSFP28 front panel uplink ports (ports 48 through 53)
  - QFX5120-32C: Any network ports installed with either 40-Gigabit Ethernet QSFP+ or 100-Gpbs QSFP28 transceivers
  - QFX5120-48YM: Only the eight QSFP+ front panel uplink port (ports 48 through 55)
- On QFX5110 switches: Any 40-Gigabit Ethernet QSFP+ or 100-Gigabit Ethernet QSFP28 ports
- On QFX3500, QFX3600, or QFX5100 switches: Any non-channelized 40-Gigabit Ethernet QSFP+ ports
- Any fixed 10-Gigabit Ethernet SFP+ ports on any QFX Series switches that support these ports

EX4650 and QFX Series switches don't have any dedicated VCPs (ports that can only be used as VCPs) or default-configured VCPs (ports that are configured as VCPs in the default factory configuration). See "Virtual Chassis Port Options" on page 41 for details on which ports on different QFX Series switches can be VCPs.

NOTE: By default, QFX5120-48YM switches come up in the default/HiGig mode and the Virtual Chassis ports come up as HiGig ports. To change this to HGoE mode, enter the `request virtual-chassis mode hgoe` command and reboot the switch, using the `request system reboot` command. To change this back to default/HiGig mode, enter the `request virtual-chassis mode hgoe disable` command, and reboot the switch using the `request system reboot` command.

You can increase VCP bandwidth between member switches by connecting multiple VCP links between the switches. When multiple VCP links interconnect the same two member switches, the links automatically form a Link Aggregation Group (LAG) bundle if they have the same speeds. For example, if you have two 40-Gigabit Ethernet QSFP+ interfaces configured as VCPs between member switches, the two links form a LAG with 80-Gbps of total bandwidth. However, 10-Gigabit Ethernet SFP+ and 40-Gigabit Ethernet QSFP+ interfaces configured as VCPs between two Virtual Chassis member switches will not become member links in the same VCP LAG.

You can set up a Virtual Chassis using either a nonprovisioned or a preprovisioned configuration. If you want to deterministically control the role and member ID assigned to each member switch, use a preprovisioned configuration. Virtual Chassis member switches can have one of three roles: primary Routing Engine, backup Routing Engine, or linecard role. In some combinations of switches in a Virtual Chassis, we recommend or require that you configure certain switches into the Routing Engine roles. See "Understanding Virtual Chassis Components" on page 35 for more information about Virtual Chassis roles.

Adding switches to a preprovisioned configuration is simpler if you use the automatic VCP conversion feature (see "Automatic Virtual Chassis Port (VCP) Conversion" on page 45), which automatically converts uplink ports into VCPs on the member switches on both sides of the new VCP links as they are cabled. This method to expand a Virtual Chassis is also called *autoprovisioning*.

QFX5200 Switches in a Virtual Chassis

Virtual Chassis is supported on QFX5200 switches starting in Junos OS Release 17.3R2 and 17.4R1 onward.

You can interconnect up to three QFX5200 switches into a QFX5200 Virtual Chassis. QFX5200 switches cannot be combined with other types of switches in a Virtual Chassis.

The following QFX5200 switches are supported in a QFX5200 Virtual Chassis:

- QFX5200-32C

QFX5120 or EX4650 Switches in a Virtual Chassis

QFX5120 and EX4650 switches are similar and operate the same way in a Virtual Chassis. Each QFX5120 or EX4650 switch model can only be combined with the same model of switches into a Virtual Chassis.

- Starting in Junos OS Release 19.3R1, you can interconnect two QFX5120-48Y switches or two EX4650-48Y switches into a Virtual Chassis.

  Starting in Junos OS Release 20.2R1, you can interconnect two QFX5120-48T switches into a Virtual Chassis.

  Starting in Junos OS Release 20.3R1, you can interconnect two QFX5120-32C switches into a Virtual Chassis.

  The two member switches must be in the primary and backup Routing Engine roles.

- Starting in Junos OS Release 20.1R1, you can interconnect up to four EX4650-48Y switches into a Virtual Chassis.

  We recommend that you configure two member switches in the primary and backup Routing Engine roles for Routing Engine redundancy, and the remaining switches in linecard role.

See "Understanding Virtual Chassis Components" on page 35 for more information about Virtual Chassis member roles.

QFX5120 and EX4650 switches do not have dedicated or default-configured VCPs, but you can set any of the following ports as VCPs:

- On QFX5120-48Y, QFX5120-48T, or EX4650-48Y switches: Any of the 40-Gigabit Ethernet QSFP+ or 100-Gigabit Ethernet QSFP28 uplink ports on the front panel (non-channelized). These are ports 48 through 55 on EX4650-48Y or QFX5120-48Y switches, and ports 48 through 53 on QFX5120-48T switches.

  > **NOTE**: You can't use any of the other ports (network ports 0 through 47) as VCPs. Running the `request virtual-chassis vc-port set` command on the network ports doesn't fail, but they will not function properly as VCPs.

- On QFX5120-32C switches: Any of the 32 network ports installed with either 40-Gigabit Ethernet QSFP+ or 100-Gigabit Ethernet QSFP28 transceivers

You configure, monitor, and maintain a QFX5120 Virtual Chassis or an EX4650 Virtual Chassis in the same way as other QFX Series Virtual Chassis. See the following for more details:

- "Configuring an EX4650 or a QFX Series Virtual Chassis" on page 100
- "Adding a New Switch to an Existing EX4650 or QFX Series Virtual Chassis" on page 116
- *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*

QFX5110 Switches in a Virtual Chassis

Starting in Junos OS Release 17.3R1, QFX5110 switches support Virtual Chassis.

You can interconnect up to ten QFX5110 switches or a combination of QFX5110 and QFX5100 switches into a QFX5110 Virtual Chassis. All switches can run the same software image, and you do not need to configure mixed mode.

NOTE: When using the QFX5110 Virtual Chassis, sometimes the master could reboot and lose connection with dual-homed video streams along with issues related to vCMTS device's connectivity for about two to three minutes. This is because, each daemon including the rpd daemon, is terminated consecutively by the system when a user does a manual reboot. When rpd daemons are shut down, TCP FIN packets are produced. When the FIN packet is created, two scenarios are possible. Either the interface is offline or it is still online. It is more normal for the interface to be active at this period. The BGP connection is subsequently cut off by sending a FIN message to the peer. This explains the reason why the BGP session was terminated.

The following QFX5110 and QFX5100 switches are supported in a QFX5110 Virtual Chassis:

- QFX5110-32Q
- QFX5110-48S
- QFX5100-24Q
- QFX5100-48S
- QFX5100-48T

  Starting in Junos OS Release 17.3R2, you can include QFX5100-48T switches in a QFX5110 Virtual Chassis.
- QFX5100-96S

QFX5100 Switches in a Virtual Chassis

Starting in Junos OS Release 13.2X51-D20, QFX5100 switches support Virtual Chassis. You can interconnect up to ten switches into a QFX5100 Virtual Chassis that can include any combination of QFX5100, QFX3500, QFX3600, and EX4300 switches.

NOTE: In Junos OS release 13.2X51-D20, you can interconnect only up to four QFX5100-96S switches in a non-mixed QFX5100 Virtual Chassis. Starting in Junos OS release 13.2X51-D25, you can configure up to ten QFX5100-96S switches into a mixed or non-mixed QFX Series Virtual Chassis.

The following QFX5100 switches are supported in a non-mixed QFX5100 Virtual Chassis or a mixed QFX5100 Virtual Chassis (QFX5100 switches with QFX3500, QFX3600, or EX4300 switches):

- QFX5100-24Q
- QFX5100-48S
- QFX5100-48T
- QFX5100-96S

NOTE: You should configure QFX5100-24Q switches as the Routing Engine role member switches in a mixed QFX5100 Virtual Chassis.

QFX5100 switches can also be part of a QFX5110 Virtual Chassis, which can be a combination of up to ten QFX5110 and QFX5100 switches. For a QFX5100 switch to participate in a QFX5110 Virtual Chassis, the switch must have the same software version and image installed as the software running on the QFX5110 switches in the Virtual Chassis. You don't need to configure mixed mode. In a QFX5110 Virtual Chassis, we recommend to use QFX5110 switches in the primary and backup Routing Engine roles, and QFX5100 switches only in the line-card role.

The following QFX5100 switches are supported in a QFX5110 Virtual Chassis:

- QFX5100-24Q
- QFX5100-48S
- QFX5100-48T

  Starting in Junos OS Release 17.3R2, you can include QFX5100-48T switches in a QFX5110 Virtual Chassis.

- QFX5100-96S

CAUTION: You *must* first upgrade a QFX5100 switch running a Junos OS software image with "-qfx-5-" in the package filename to a Junos OS software image with "-qfx-5e-" in the package filename before it can be added to a QFX5110 Virtual Chassis or VCF. See *Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric*.

QFX3500 and QFX3600 Switches in a Virtual Chassis

QFX3500 and QFX3600 switches support Virtual Chassis. To be included in a Virtual Chassis, you must configure QFX3500 and QFX3600 switches as standalone switches and not as QFX node devices in a QFabric system.

QFX3500 and QFX3600 switches can be in a Virtual Chassis with up to ten member switches in a non-mixed Virtual Chassis as follows:

- All QFX3500 switches
- All QFX3600 switches
- A combination of QFX3500 and QFX3600 switches

QFX3500 or QFX3600 switches can also be in a mixed QFX5100 Virtual Chassis with any combination of QFX3500, QFX3600, QFX5100, and EX4300 switches.

NOTE: EX4300 multigigabit model (EX4300-48MP) switches are not supported in a mixed Virtual Chassis with QFX Series switches.

EX4300 Switches in a QFX Series Virtual Chassis

Starting in Junos OS Release 13.2X51-D20, EX4300 switches except multigigabit models (EX4300-48MP) can be interconnected into a mixed-mode QFX Series Virtual Chassis with up to ten member switches that can be any combination of EX4300, QFX3500 switches, QFX3600 switches, and QFX5100 switches.

Release History Table

Release	Description
20.2R1	Starting in Junos OS Release 20.2R1, you can interconnect two QFX5120-48T switches into a Virtual Chassis.
20.1R1	Starting in Junos OS Release 20.1R1, you can interconnect up to four EX4650-48Y switches into a Virtual Chassis.
19.4R1	Starting in Junos OS Release 20.3R1, you can interconnect two QFX5120-32C switches into a Virtual Chassis.

19.3R1	Starting in Junos OS Release 19.3R1, you can interconnect two QFX5120-48Y or EX4650-48Y switches into a Virtual Chassis.
17.3R2-S4	Starting in Junos OS Release 17.3R2-S4, 100-Gigabit Ethernet QSFP28 ports are also supported as VCPs on QFX5200 switches.
17.3R2	Virtual Chassis is supported on QFX5200 switches starting in Junos OS Release 17.3R2 and 17.4R1 onward.
17.3R2	Starting in Junos OS Release 17.3R2, you can include QFX5100-48T switches in a QFX5110 Virtual Chassis.
17.3R1	Starting in Junos OS Release 17.3R1, QFX5110 switches support Virtual Chassis.
13.2X53-D25	Starting in Junos OS release 13.2X51-D25, you can configure up to ten QFX5100-96S switches into a mixed or non-mixed QFX Series Virtual Chassis.
13.2X51-D20	Starting in Junos OS Release 13.2X51-D20, QFX5100 switches support Virtual Chassis.
13.2X51-D20	In Junos OS release 13.2X51-D20, you can interconnect only up to four QFX5100-96S switches in a non-mixed QFX5100 Virtual Chassis.
13.2X51-D20	Starting in Junos OS Release 13.2X51-D20, EX4300 switches except multigigabit models (EX4300-48MP) can be interconnected into a mixed-mode QFX Series Virtual Chassis with up to ten member switches that can be any combination of EX4300, QFX3500 switches, QFX3600 switches, and QFX5100 switches.

RELATED DOCUMENTATION

Virtual Chassis Overview for Switches | 2
Understanding Virtual Chassis Components | 35
Configuring an EX4650 or a QFX Series Virtual Chassis | 100

Understanding Virtual Chassis Components

IN THIS SECTION

This topic describes the components of an EX series or a QFX Series *Virtual Chassis*.

- An EX Series Virtual Chassis is a supported combination of standalone EX Series switches interconnected and managed as a single chassis. This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis.

  See *Understanding EX8200 Virtual Chassis Components* for information about EX8200 Virtual Chassis.

  > **NOTE**: We do not recommend using EX9200 switches in a Virtual Chassis, and we phased out support for that architecture as of Junos OS Release 17.1R1. For deployments with EX9200 switches, we recommend planning or moving to MC-LAG or Junos Fusion Enterprise architectures instead of using a Virtual Chassis.

- A QFX Series Virtual Chassis is a supported combination of standalone QFX3500, QFX3600, QFX5100, QFX5110, QFX5120, or QFX5200 switches interconnected and managed as a single chassis. EX4650 Virtual Chassis operate the same as QFX5120 Virtual Chassis, so most of the information in this topic about QFX Series Virtual Chassis in general also applies to an EX4650 Virtual Chassis, with a few platform-specific support differences.

NOTE: EX4300 switches (excluding multigigabit models [EX4300-48MP]) can also be interconnected into a mixed Virtual Chassis with QFX3500, QFX3600, and QFX5100 switches.

This topic does not describe Virtual Chassis Fabric components. Instead, see *Understanding Virtual Chassis Fabric Components*.

Maximum Switch Support

The maximum number of switches that a Virtual Chassis supports varies by Virtual Chassis and might also depend on the Junos OS release running on the Virtual Chassis.

Maximum Number of Switches in an EX Series Virtual Chassis

Table 3 on page 36 lists the maximum number of member switches supported in an EX Series Virtual Chassis by Junos OS release.

Table 3: Maximum Member Switch Support for EX Series Virtual Chassis by Junos OS Release

Type of EX Series Virtual Chassis	Maximum Member Switches by Junos OS Release
EX2200 Virtual Chassis	12.2R1—Initial release. Up to 4 EX2200 member switches.
EX2300 Virtual Chassis	15.1X53-D50—Initial release. Up to 4 EX2300 member switches. 18.1R2—Up to 4 multigigabit EX2300 (EX2300-24MP and EX2300-48MP) member switches. 18.4R1—Starting in Junos OS Release 18.4R1, up to 4 of any model EX2300 member switches (including multigigabit models and any other EX2300 switches) can be combined in the same Virtual Chassis.

Table 3: Maximum Member Switch Support for EX Series Virtual Chassis by Junos OS Release *(Continued)*

Type of EX Series Virtual Chassis	Maximum Member Switches by Junos OS Release
EX3300 Virtual Chassis	11.3R1—Initial release. Up to 6 EX3300 member switches. 12.2R1—Starting in Junos OS Release 12.2R1, an EX3300 Virtual Chassis can support up to 10 EX3300 member switches.
EX3400 Virtual Chassis	15.1X53-D50—Initial release. Up to 10 EX3400 member switches.
EX4100 Virtual Chassis	22.2R1—Initial release. Up to 10 EX4100 member switches.
EX4200 Virtual Chassis	9.0R1—Initial release. Up to 10 EX4200 member switches.
EX4300 Virtual Chassis	13.2X50-D10—Initial release. Up to 10 EX4300 member switches. 13.2X50-D20—Starting in Junos OS Release 13.2X50-D20, a mixed QFX Series Virtual Chassis or VCF can also contain EX4300 switches. 18.2R1—Starting in Junos OS Release 18.2R1 with the introduction of EX4300 multigigabit model switches (EX4300-48MP), an EX4300 Virtual Chassis can contain up to 10 EX4300 multigigabit model switches as a non-mixed Virtual Chassis or a combination of EX4300 multigigabit model switches with other EX4300 switches as a mixed EX4300 Virtual Chassis.
EX4400 Virtual Chassis	21.1R1—Initial release. Up to 10 EX4400 member switches. 21.2R1—Starting in Junos OS Release 21.2R1, an EX4400 Virtual Chassis can also include EX4400 multigigabit model switches (EX4400-24MP and EX4400-48MP).

Table 3: Maximum Member Switch Support for EX Series Virtual Chassis by Junos OS Release *(Continued)*

Type of EX Series Virtual Chassis	Maximum Member Switches by Junos OS Release
EX4500 Virtual Chassis	11.1R1—Initial release. Support for up to 2 EX4500 member switches. 11.4R1—Support for up to 10 EX4500 member switches.
EX4550 Virtual Chassis	12.2R1—Initial release. Support for up to 10 EX4550 member switches.
EX4600 Virtual Chassis	13.2X51-D25—Initial release. Support for up to 10 EX4600 member switches.
EX4650 Virtual Chassis	19.3R1—Initial release. Up to 2 EX4650 switches in Routing Engine roles only. 20.1R1—Starting in Junos OS Release 20.1R1, an EX4650 Virtual Chassis can have up to 4 members.
Mixed EX4200 and EX4500 Virtual Chassis	11.1R1—Initial release. Up to 2 EX4500 member switches and up to 8 EX4200 member switches. 11.2R1—Up to 9 EX4200 member switches. 11.4R1—Up to 9 EX4500 member switches.
Mixed EX4200 and EX4550 Virtual Chassis	12.2R1—Initial release. Up to 10 total EX4200 and EX4550 member switches.
Mixed EX4200, EX4500, and EX4550 Virtual Chassis	12.2R1—Initial release. Up to 10 total EX4200, EX4500, and EX4550 member switches.

Table 3: Maximum Member Switch Support for EX Series Virtual Chassis by Junos OS Release *(Continued)*

Type of EX Series Virtual Chassis	Maximum Member Switches by Junos OS Release
Mixed EX4300 and EX4600 Virtual Chassis	13.2X51-D25—Initial release. Up to 10 total EX4300 and EX4600 member switches. EX4600 member witches must assume Routing Engine role. **NOTE**: EX4300 multigigabit model (EX4300-48MP) switches are not supported in a mixed Virtual Chassis with EX4600 switches.
Mixed EX4500 and EX4550 Virtual Chassis	12.2R1—Initial release. Up to 10 total EX4500 and EX4550 switches.
EX9200 Virtual Chassis	13.2R2—Initial release. Up to 2 EX9200 switches. **NOTE**: We have phased out support for EX9200 switches in a Virtual Chassis as of Junos OS Release 17.1R1. For deployments with EX9200 switches, we recommend planning or moving to MC-LAG or Junos Fusion Enterprise architectures instead of using a Virtual Chassis configuration.

Maximum Number of Switches in a QFX Series Virtual Chassis (Including Mixed Virtual Chassis with EX Series Switches)

Table 4 on page 39 lists the maximum number of member switches supported in a QFX Series Virtual Chassis by Junos OS release, including mixed QFX Series Virtual Chassis with EX Series switch members.

Table 4: Maximum Member Switch Support for QFX Series Virtual Chassis by Junos OS Release

Type of QFX Series Virtual Chassis	Maximum Member Switches by Junos OS Release
QFX3500 or QFX3600 Virtual Chassis: • Only QFX3500 and QFX3600 switches in any combination.	13.2X50-D15—Initial release. Up to 10 total member switches.

Table 4: Maximum Member Switch Support for QFX Series Virtual Chassis by Junos OS Release *(Continued)*

Type of QFX Series Virtual Chassis	Maximum Member Switches by Junos OS Release
QFX3500 or QFX3600 mixed-mode Virtual Chassis: • QFX3500 and QFX3600 switches in Routing Engine role with any combination of QFX3500, QFX3600, and EX4300 switches (excluding EX4300 multigigabit models) in linecard role.	13.2X51-D20—Up to 10 total member switches.
QFX5100 Virtual Chassis: • Only QFX5100 switches.	13.2X51-D20—Initial release. Up to 10 member switches (except QFX5100-96S). 13.2X51-D20—Initial release for QFX5100-96S. Up to 4 member switches. 13.2X53-D25—Prior to Junos OS Release 13.2X51-D25, only up to 4 member switches can be in a QFX5100 Virtual Chassis. Starting in Junos OS Release 13.2X51-D25, you can include up to 10 QFX5100-96S switches in a mixed or non-mixed QFX5100 Virtual Chassis.
QFX5100 mixed-mode Virtual Chassis: • QFX5100 switches in Routing Engine role with any combination of QFX5100, QFX3500, QFX3600, and EX4300 switches (excluding EX4300 multigigabit models) in linecard role.	13.2X51-D20—Initial release. Up to 10 total member switches (except QFX5100-96S). 13.2X53-D25—Up to 10 total member switches (including QFX5100-96S).
QFX5110 Virtual Chassis: • QFX5110 switches in Routing Engine role with any combination of supported QFX5110 and QFX5100 switches in linecard role.	17.3R1—Initial release. Up to 10 member switches.

Table 4: Maximum Member Switch Support for QFX Series Virtual Chassis by Junos OS Release *(Continued)*

Type of QFX Series Virtual Chassis	Maximum Member Switches by Junos OS Release
QFX5120 Virtual Chassis:	19.3R1—Initial release on QFX5120-48Y switches. Up to 2 member switches, both in Routing Engine role. 20.2R1—Initial release on QFX5120-48T switches. Up to 2 member switches, both in Routing Engine role. 20.3R1—Initial release on QFX5120-32C switches. Up to 2 member switches, both in Routing Engine role.
QFX5200 Virtual Chassis— • Only QFX5200 switches.	17.3R2 and 17.4R1—Initial release. Up to 3 member switches.

Virtual Chassis Ports (VCPs)

You set up a Virtual Chassis by configuring Virtual Chassis ports (VCPs) on the member switches, and interconnecting the switches using the VCPs. VCPs are responsible for passing all data and control traffic between member switches in the Virtual Chassis.

Virtual Chassis Port Options

Some switches have dedicated VCPs; you can only use these ports as VCPs and you can’t reconfigure them as network ports. Dedicated VCPs allow you to interconnect switches into a Virtual Chassis without requiring any additional interface configuration.

Some switches have ports that are configured as VCPs by default. You don’t need to explicitly configure those as VCPs to use them to interconnect the switches into a Virtual Chassis.

Most switches have optical or uplink ports that you can also configure as VCPs.

You must configure VCPs to interconnect switches that do not have dedicated or default-configured VCPs or to interconnect switches across greater distances than allowed by a dedicated VCP connection. Otherwise, you can mix any of the supported VCP options among the members of a Virtual Chassis, and we recommend having redundant links between any two members for resiliency or to increase member communication bandwidth. VCPs automatically bundle into a Link Aggregation Group when two or more ports operating at the same speed are configured into VCPs between the same two member switches. See "Understanding Virtual Chassis Port Link Aggregation" on page 68 for details.

When adding switches to an existing Virtual Chassis or adding new redundant links between existing members, if the automatic VCP conversion feature is enabled, under the right conditions the ports on both sides of the connection will convert into VCPs automatically (see "Automatic Virtual Chassis Port (VCP) Conversion" on page 45).

Table 5 on page 42 summarizes the available VCP options on switches in an EX Series or QFX Series Virtual Chassis. For complete details on where dedicated VCPs, default-configured VCPs, or ports that can be configured as VCPs are located on a switch, and the supported transceivers and cables that you can use for VCP connections on the switch, see the hardware documentation for that type of switch.

Table 5: VCP Options by Switch Type

Switch	Dedicated VCPs	Default VCPs	Ports that can be configured and are supported as VCPs
EX2200	None	None	Any uplink ports All RJ-45 interfaces, including built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors and 1000BASE-T RJ-45 transceivers
EX2300 (including EX2300 multigigabit models)	None	None	10-Gigabit Ethernet uplink ports with SFP+ tranceivers **NOTE**: You cannot use ports with SFP transceivers as VCPs on EX2300 switches to form a Virtual Chassis.
EX3300	None	Uplink ports 2 and 3	Any of the 4 uplink ports (ports 0 through 3)
EX3400	None	All QSFP+ uplink ports (PIC slot 1, ports 0 and 1)	Any SFP+ uplink ports **NOTE**: You cannot use ports with SFP transceivers as VCPs on EX3400 switches to form a Virtual Chassis.
EX4100	4 25-Gbps SFP28 ports on the front panel	4 25-Gbps SFP28 ports on the front panel	None

Table 5: VCP Options by Switch Type *(Continued)*

Switch	Dedicated VCPs	Default VCPs	Ports that can be configured and are supported as VCPs
EX4200	2 ports on the rear panel	None	Any uplink module ports (SFP, SFP+, or XFP) or through an SFP+ port on the EX4200-24F switch **NOTE**: You cannot set a 1000BASE-T copper SFP transceiver (EX-SFP-1GE-T) connection as a VCP on EX4200 switches.
EX4300	None	All QSFP+ ports	Any uplink ports installed with SFP+ or QSPF+ transceivers **NOTE**: On 32-port EX4300 switches, you can't use the four built-in 10-Gigabit Ethernet SFP+ ports as VCPs.
EX4300 multigigabit models (EX4300-48MP)	4 40-Gbps QSFP+ ports on the rear panel	None	None
EX4400 (Including EX4400 multigigabit models)	None	4 logical 50-Gbps VCP interfaces using the two 100-Gbps QSFP28 ports on the rear panel (PIC slot 1)	None
EX4500 and EX4550	Two ports on the Virtual Chassis module	None	Any SFP, SFP+, or XFP uplink module port **NOTE**: You cannot use SFP+ uplink ports installed with 1000BASE-T copper SFP transceivers (EX-SFP-1GE-T) as VCP connections on EX4500 and EX4550 switches.
EX4600	None	None	Any SFP+ and QSFP+ ports

Table 5: VCP Options by Switch Type *(Continued)*

Switch	Dedicated VCPs	Default VCPs	Ports that can be configured and are supported as VCPs
EX4650	None	None	Any of the 40-Gigabit Ethernet or 100-Gigabit QSFP28 ports on the front panel (ports 48 through 55), non-channelized **NOTE**: The Junos OS doesn’t prevent you from trying to set other ports as VCPs, but they don’t operate properly as VCPs.
QFX3500 and QFX3600	None	None	Any non-channelized 40-Gigabit Ethernet QSFP + interfaces
QFX5100	None	None	Any non-channelized 40-Gigabit Ethernet QSFP + interfaces
QFX5110	None	None	Any 40-Gigabit Ethernet or 100-Gigabit Ethernet QSFP28 ports Any non-channelized 40-Gigabit Ethernet QSFP + interfaces Any non-channelized 10-Gigabit Ethernet SFP+ interfaces (on QFX5110 switch models that support these ports)

Table 5: VCP Options by Switch Type *(Continued)*

Switch	Dedicated VCPs	Default VCPs	Ports that can be configured and are supported as VCPs
QFX5120	None	None	(QFX5120-48Y) Any of the eight 40-Gigabit Ethernet or 100-Gigabit Ethernet QSFP+ or QSFP28 ports on the front panel (ports 48 through 55), non-channelized (QFX5120-48T) Any of the six 40-Gigabit Ethernet or 100-Gigabit Ethernet QSFP+ or QSFP28 ports on the front panel (ports 48 through 53), non-channelized **NOTE**: Any ports other than those specified above for QFX5120-48Y and QFX5120-48T switches are not supported as VCPs. The Junos OS CLI doesn’t return an error if you try to set other ports as VCPs, but they will not work properly as VCPs. (QFX5120-32C) Any non-channelized network ports (ports 0 through 31) installed with either 40-Gbps QSFP+ or 100-Gpbs QSFP28 transceivers
QFX5200	None	None	Any 40-Gigabit Ethernet QSFP+ ports Starting in Junos OS Release 17.3R2-S4, you can also use 100-Gigabit Ethernet QSFP28 ports as VCPs on QFX5200 switches.

All supported SFP, SFP+, and XFP uplink connections between EX4200, EX4500, and EX4550 switches can be configured as VCPs.

QSFP+ interfaces that have been channelized into SFP+ interfaces using a breakout cable cannot be configured into VCPs.

Automatic Virtual Chassis Port (VCP) Conversion

When the automatic VCP conversion feature is enabled and you cable a new link from a new switch being added into an existing Virtual Chassis, or add a redundant link between two members of a Virtual Chassis, ports that can be VCPs are automatically converted into VCPs under the following conditions:

- Link Layer Discovery Protocol (LLDP) or LLDP-Media Endpoint Discovery (LLDP-MED) is enabled on the interfaces for the members on both ends of the new link. The two sides exchange LLDP packets to accomplish the port conversion.

- The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the `set virtual-chassis member` command.

- The interfaces for the ports on both ends of the link are not already configured as VCPs. Both sides of the link must be in the same state to handshake and establish the VCP link.

Using automatic VCP conversion when adding a switch to a preprovisioned Virtual Chassis is also called *autoprovisioning* the new member.

For ports to be eligible for automatic VCP conversion, you must convert them back into network ports using the `request virtual-chassis vc-port delete` command if they are default-configured VCPs or you previously configured them into VCPs. Switches do not automatically convert VCPs back into network ports when you remove them from a Virtual Chassis and disconnect the links.

Automatic VCP conversion is enabled by default on all Virtual Chassis, except in the following cases:

- Automatic VCP conversion doesn't apply to EX4400 switches in a Virtual Chassis. On these switches, to convert the default VCPs into network ports or convert them from network ports back into VCP ports, you must explicitly set the port mode using the `request virtual-chassis mode network-port` command, and then reboot the switch.

- Starting in Junos OS Releases 15.1R7 and 14.1X53-D47, in EX2200, EX3300, EX4200, EX4500, and EX4550 Virtual Chassis, automatic VCP conversion is disabled by default. If desired, you can enable the feature by configuring the `auto-conversion` statement at the `[edit virtual-chassis]` hierarchy level on the Virtual Chassis.

  > **CAUTION**: When automatic VCP conversion is enabled in a Virtual Chassis with switches that have dedicated VCPs (EX4200, EX4500, or EX4550 Virtual Chassis), if network or uplink ports are automatically converted into VCPs to create a redundant link with a dedicated VCP connection between the same two Virtual Chassis members, you must reboot the Virtual Chassis to avoid creating a traffic loop within the Virtual Chassis. (The same issue can happen even if you manually convert the ports into VCPs to create the redundant VCP link with a dedicated VCP link, so you must also reboot the Virtual Chassis in that case to avoid traffic looping.)

- Starting in Junos OS Releases 14.1X53-D47, 17.4R2, 18.1R3, 18.2R2, and 18.3R1 for EX4300, EX4600, QFX Series Virtual Chassis and for any EX4650 and QFX5120 Virtual Chassis (which all have the automatic VCP conversion feature enabled by default), you can choose to disable the feature by configuring `no-auto-conversion` at the `[edit virtual-chassis]` hierarchy level on the Virtual

Chassis. To return to the default behavior to re-enable automatic VCP conversion, delete the `no-auto-conversion` statement from the configuration.

Virtual Chassis Port Link Aggregation Groups

You can increase VCP bandwidth between member switches by configuring multiple links between the same two switches into VCP links. When multiple VCPs interconnect the same two member switches, the links automatically form a Link Aggregation Group (LAG) bundle if the VCP links are the same speed. For example, if you have two 40-Gbps QSFP+ VCP links connected between member switches, the links automatically form a LAG with 80-Gbps total bandwidth. However, 10-Gigabit SFP+ and 40-Gbps QSFP + VCP links will not become members of the same LAG.

Within a Virtual Chassis, you can also configure network interfaces located on different Virtual Chassis member switches to form a LAG, which provides load-balancing and redundancy for network traffic that the Virtual Chassis forwards. See "Understanding Virtual Chassis Port Link Aggregation" on page 68 for details on the difference between VCP LAGs and network interface LAGs within a Virtual Chassis.

Primary Routing Engine Role

In a Virtual Chassis, each member switch operates in one of two roles, Routing Engine role or linecard role. When in Routing Engine role, a member switch acts as the primary or backup Routing Engine.

The primary Routing Engine member in the Virtual Chassis:

- Manages the member switches.
- Runs Junos OS for the switches as a primary Routing Engine.
- Runs the chassis management processes and control protocols.
- Represents all the member switches interconnected within the Virtual Chassis configuration. (The hostname and other properties that you assign to this switch during setup apply to all members of the Virtual Chassis configuration.)

In a preprovisioned configuration, the Virtual Chassis primary-role election algorithm determines which member switch in the Routing Engine role acts as the Virtual Chassis primary and which acts as the backup. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62.

In a configuration that is not preprovisioned, called a *nonprovisioned* configuration, the Virtual Chassis selects the primary and backup using the primary-role priority value and secondary factors in the primary-role election algorithm.

The remaining switches in the Virtual Chassis that are not the primary or backup operate in the linecard role.

Use the following guidelines for assigning Routing Engine roles to the switches in a mixed Virtual Chassis:

- In any mixed Virtual Chassis configuration that includes EX4200 switches, EX4500 switches, or EX4550 switches, you can configure any switch in any role in any configuration.
- In a mixed EX4300 Virtual Chassis composed of EX4300 multigigabit model (EX4300-48MP) and other EX4300 model switches, you should always have EX4300 multigigabit model switches in the Routing Engine role.
- In a mixed EX4600 Virtual Chassis with EX4300 switches, EX4600 switches are always in the primary Routing Engine role.
- In a mixed QFX Series Virtual Chassis made up of QFX5100 switches with QFX3500, QFX3600, or EX4300 switches, we recommend configuring QFX5100 switches into the Routing Engine role.
- In a mixed QFX Series Virtual Chassis made up of QFX3500, QFX3600, and EX4300 switches, we recommend configuring QFX3500 or QFX3600 switches into the primary and backup Routing Engine roles.
- In a QFX5110 Virtual Chassis with QFX5110 and QFX5100 switches, we recommend configuring only QFX5110 switches into the Routing Engine role.
- In a two-member EX4650 or QFX5120 Virtual Chassis, configure both member switches into the Routing Engine role as primary and backup member switches only (no linecard role members).

Backup Routing Engine Role

The member that functions in the backup Routing Engine role in a Virtual Chassis:

- Maintains a state of readiness to take over the primary Routing Engine role if the primary fails.
- Runs Junos OS for the switches as a backup Routing Engine.
- Synchronizes with the primary in terms of protocol states, forwarding tables, and other information, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the primary is unavailable.

The Virtual Chassis configuration must have at least two member switches in order to have a backup Routing Engine member.

In a preprovisioned configuration, the Virtual Chassis primary-role election algorithm determines which member switch in the Routing Engine role acts as the Virtual Chassis primary and which acts as the backup. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62.

In a nonprovisioned configuration, the Virtual Chassis selects the primary and backup member switches using the primary-role priority value and secondary factors in the primary-role election algorithm.

Use the following guidelines for assigning Routing Engine roles to the switches in a mixed Virtual Chassis:

- In any mixed Virtual Chassis configuration that includes EX4200 switches, EX4500 switches, or EX4550 switches, you can configure any switch in any role in any configuration.
- In a mixed EX4300 Virtual Chassis composed of EX4300 multigigabit model (EX4300-48MP) and other EX4300 model switches, you should always have EX4300 multigigabit model switches in the primary and backup Routing Engine roles.
- In a mixed EX4600 Virtual Chassis with EX4300 switches, you must use an EX4600 switch in the primary role, and we strongly recommend you also configure an EX4600 switch into the backup role to help the Virtual Chassis remain stable when a routing engine switchover happens.
- In a mixed QFX Series Virtual Chassis made up of QFX5100 switches with QFX3500, QFX3600, or EX4300 switches, we recommend configuring the QFX5100 switches into the Routing Engine role.
- In a mixed QFX Series Virtual Chassis made up of QFX3500, QFX3600, and EX4300 switches, we recommend configuring only QFX3500 or QFX3600 switches into the primary and backup Routing Engine roles.
- In a QFX5110 Virtual Chassis with QFX5110 and QFX5100 switches, we recommend configuring only QFX5110 switches into the Routing Engine role.
- In a two-member EX4650 or QFX5120 Virtual Chassis, configure both member switches into the Routing Engine role as primary and backup member switches only (no linecard role members).

Linecard Role

A member that functions in the linecard role in a Virtual Chassis:

- Runs only a subset of Junos OS.
- Does not run the chassis control protocols.
- Can detect certain error conditions (such as an unplugged cable) on any interfaces that have been configured on it through the primary.

The Virtual Chassis configuration must have at least three members in order to include a linecard member.

In a preprovisioned configuration, you can explicitly configure a member with the linecard role, which means it can't be a primary or backup Routing Engine.

In a nonprovisioned configuration, the members that are not selected as primary or backup operate as linecard members of the Virtual Chassis. The Virtual Chassis selects the primary and backup member switches using the primary-role priority value and secondary factors in the primary-role election algorithm. A switch with a primary-role priority of 0 is always in the linecard role.

In any two-member Virtual Chassis, for high availability you should configure both members into the Routing Engine role, and no members in the linecard role. Otherwise, in a Virtual Chassis with more than two members, any supported switch type can operate in linecard role.

Use the following guidelines for assigning Routing Engine and linecard roles to the switches in a QFX Series Virtual Chassis:

- In a mixed QFX Series Virtual Chassis made up of QFX5100 switches with QFX3500, QFX3600, or EX4300 switches, we recommend configuring the QFX5100 switches into the Routing Engine role. In a mixed QFX Series Virtual Chassis that does not contain QFX5100 switches, we recommend configuring QFX3500 or QFX3600 switches into the Routing Engine role.

- In a QFX5110 Virtual Chassis made up of QFX5110 and QFX5100 switches, we recommend configuring only QFX5110 switches into the Routing Engine role.

Member Switch and Member ID

Each standalone switch that supports Virtual Chassis is a potential member of a Virtual Chassis configuration. When you power on one of those switches, it has a Virtual Chassis member ID that you can see on the front-panel LCD on some switches or in `show virtual-chassis` command output. If the switch is powered on as a standalone switch, its member ID is always `0`. When you interconnect the switch into a Virtual Chassis configuration, the primary member switch assigns it a member ID based on various factors such as the order in which the switch was added to the Virtual Chassis or if you defined member IDs based on switch serial numbers in the preprovisioning process.

If the Virtual Chassis configuration previously included a member switch and you physically disconnected or removed that member from the Virtual Chassis configuration, its member ID is not automatically available for assignment as part of the primary's standard sequential member ID assignment. For example, you might have a Virtual Chassis configuration with member 0, member 2, and member 3, because member 1 was removed. When you add another member switch and power it on, the primary assigns ID 4 to it, not ID 1. If you want to reuse a member ID from a member switch that was removed, you can *recycle* the member id (see the `request virtual-chassis recycle` command for details). .

The member ID distinguishes the member switches from each other. You use the member ID to:

- assign a primary-role priority value to a member switch.
- configure interfaces for a member switch, similar to specifying a juniper Networks device slot number.
- apply some operational commands to a member switch.
- display status or characteristics of a member switch.

Primary-role Priority

In a nonprovisioned configuration, you can designate the role (primary or backup Routing Engine role or linecard role) that a member switch assumes by configuring its primary-role priority (a number from `0` through `255`). The primary-role priority value is the first consideration in the primary-role election algorithm for selecting the primary of the Virtual Chassis configuration. A switch with a primary-role priority of `0` never assumes the backup or primary Routing Engine role.

When you power on a standalone switch, it has the default primary-role priority value `128`. Because it’s the only member switch in its own Virtual Chassis configuration, it’s also the primary member. When you interconnect a standalone switch to an existing Virtual Chassis configuration (which already has its own primary), we recommend that you explicitly configure the primary-role priority of the members that you want to function as the primary and backup.

> **NOTE**: Configuring the same primary-role priority value for both the primary and backup helps to ensure a smooth transition from primary to backup if the primary becomes unavailable. It prevents the original primary from preempting control from the backup when the backup has taken control of the Virtual Chassis configuration because the original primary became unavailable.

In a preprovisioned configuration, you can’t configure primary-role priority values manually. You assign the role of each member switch, and the Virtual Chassis assigns the primary-role priority automatically based on the assigned role.

Virtual Chassis Identifier (VCID)

All members of a Virtual Chassis configuration share one Virtual Chassis identifier (VCID). The Virtual Chassis derives this identifier from internal parameters. When you monitor a Virtual Chassis configuration, certain interface views and the `show virtual-chassis` command display the VCID.

Nonvolatile Storage in a Virtual Chassis

EX Series and QFX Series switches store Junos OS system files in internal flash memory. In Virtual Chassis configurations, both the primary and the backup switch store the configuration information for all the member switches.

Junos OS optimizes the way a Virtual Chassis stores its configuration if a member switch or the Virtual Chassis configuration shuts down improperly, as follows:

- If the primary is not available, the backup switch takes on the role of the primary and its internal flash memory takes over as the alternate location for maintaining nonvolatile configuration memory.
- If you take a member switch offline for repair, the primary stores the configuration of the member switch.

NOTE: File storage management differs in an EX8200 Virtual Chassis; see *Understanding File Storage in an EX8200 Virtual Chassis* for details.

Release History Table

Release	Description
20.1R1	Starting in Junos OS Release 20.1R1, an EX4650 Virtual Chassis can have up to 4 members.
18.4R1	Starting in Junos OS Release 18.4R1, up to 4 of any model EX2300 member switches (including multigigabit models and any other EX2300 switches) can be combined in the same Virtual Chassis.
18.2R1	Starting in Junos OS Release 18.2R1 with the introduction of EX4300 multigigabit model switches (EX4300-48MP), an EX4300 Virtual Chassis can contain up to 10 EX4300 multigigabit model switches as a non-mixed Virtual Chassis or a combination of EX4300 multigigabit model switches with other EX4300 switches as a mixed EX4300 Virtual Chassis.
17.3R2-S4	Starting in Junos OS Release 17.3R2-S4, you can also use 100-Gigabit Ethernet QSFP28 ports as VCPs on QFX5200 switches.
15.1R7	Starting in Junos OS Releases 15.1R7 and 14.1X53-D47, in EX2200, EX3300, EX4200, EX4500, and EX4550 Virtual Chassis, automatic VCP conversion is disabled by default.

14.1X53-D47	Starting in Junos OS Releases 14.1X53-D47, 17.4R2, 18.1R3, 18.2R2, and 18.3R1 for EX4300, EX4600, QFX Series Virtual Chassis and for any EX4650 and QFX5120 Virtual Chassis (which all have the automatic VCP conversion feature enabled by default), you can choose to disable the feature by configuring `no-auto-conversion` at the `[edit virtual-chassis]` hierarchy level on the Virtual Chassis.
13.2X53-D25	Starting in Junos OS Release 13.2X51-D25, you can include up to 10 QFX5100-96S switches in a mixed or non-mixed QFX5100 Virtual Chassis.
13.2X50-D20	Starting in Junos OS Release 13.2X50-D20, a mixed QFX Series Virtual Chassis or VCF can also contain EX4300 switches.
12.2R1	Starting in Junos OS Release 12.2R1, an EX3300 Virtual Chassis can support up to 10 EX3300 member switches.

RELATED DOCUMENTATION

Virtual Chassis Overview for Switches | 2
Understanding EX8200 Virtual Chassis Components
Understanding EX Series Virtual Chassis | 9
Understanding QFX Series Virtual Chassis | 25
Understanding Mixed EX Series and QFX Series Virtual Chassis | 54
Configuring an EX4650 or a QFX Series Virtual Chassis | 100
Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port | 133
Command Forwarding Usage with EX Series and QFX Series Virtual Chassis | 147
Example: Configuring an EX4200 Virtual Chassis with a Primary and Backup in a Single Wiring Closet
Example: Configuring an EX4500 Virtual Chassis with a Primary and Backup in a Single Wiring Closet
Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File

Understanding Mixed EX Series and QFX Series Virtual Chassis

IN THIS SECTION

This topic describes the requirements for a mixed Virtual Chassis.

A *mixed Virtual Chassis* includes two or more types of EX Series switches, two or more types of QFX Series switches, or a supported combination of EX and QFX Series switches, where architectural differences require the Virtual Chassis to be configured into mixed mode for the switches to interoperate.

A Virtual Chassis composed of all the same type of switch can usually operate as a *non-mixed Virtual Chassis*, which does not require you to set mixed mode. However, the following Virtual Chassis that include different models of the same product must operate in mixed mode due to architecture differences between the different models:

- An EX4300 Virtual Chassis composed of EX4300 multigigabit model (EX4300-48MP) switches mixed with any other EX4300 model switches

  NOTE: An EX4300 Virtual Chassis operates as a non-mixed Virtual Chassis if it is composed of only EX4300 multigigabit model switches, or composed of any combination of any other EX4300 switches (excluding the multigigabit models).

The following combinations of different switch types can be interconnected into a Virtual Chassis that does *not* require you to set mixed mode because the switches can run the same software image when in a Virtual Chassis:

- An EX Series Virtual Chassis composed of only EX4500 and EX4550 switches
- An EX2300 Virtual Chassis composed of any models of EX2300 and EX2300 multigigabit switches

> **NOTE**: Junos OS releases prior to 18.4R1 support forming an EX2300 Virtual Chassis using only EX2300 multigigabit switches or only EX2300 switches that are not multigigabit model switches. Starting in Junos OS Release 18.4R1, EX2300, EX2300-C, and EX2300 multigigabit switches can all be combined in the same non-mixed Virtual Chassis.

- An EX4400 Virtual Chassis composed of any models of EX4400 and EX4400 multigigabit switches
- A QFX Series Virtual Chassis composed of only QFX3500 and QFX3600 switches
- A QFX5110 Virtual Chassis composed of QFX5110 and supported QFX5100 switches

QFX5200, QFX5120, EX4650, EX4400, EX3400, EX3300, EX2300, and EX2200 switches cannot be mixed with any other models of switches in a Virtual Chassis.

Mixed and Non-mixed EX Series and QFX Series Virtual Chassis Summary

Table 6 on page 55 provides a high-level overview of the EX Series and QFX Series switches allowed in the Routing Engine and line-card roles of supported mixed and non-mixed Virtual Chassis configurations. Any supported mixed or non-mixed combination of switches can be configured as a nonprovisioned or preprovisioned Virtual Chassis.

Table 6: Virtual Chassis Summary

Category	Allowed Routing Engine Members	Allowed Line Card Members
Non-mixed	QFX5200	QFX5200

Table 6: Virtual Chassis Summary *(Continued)*

Category	Allowed Routing Engine Members	Allowed Line Card Members
	QFX5120-48Y or QFX5120-48T or QFX5120-32C (2 switches of the same model only)	None, only 2 members of the same model switch are supported in any QFX5120 Virtual Chassis, and both must be in Routing Engine role
	QFX5110	QFX5110 QFX5100 (with a "-qfx-5e-" Junos OS image)
	QFX5100	QFX5100
	QFX3600 QFX3500	QFX3600 QFX3500
	EX4100	EX4100
	EX4650	None if only 2 members in the Virtual Chassis (both must be in Routing Engine role), or only other EX4650 switches otherwise. Starting in Junos OS Release 20.1R1, you can interconnect up to four EX4650 switches in an EX4650 Virtual Chassis. You should configure two member switches into the primary and backup Routing Engine roles, and the remaining member switches into the linecard role.
	EX4600	EX4600
	EX4400 (including multigigabit and any other models)	EX4400 (including multigigabit and any other models)
	EX4300 multigigabit models (EX4300-48MP) only	EX4300 multigigabit models (EX4300-48MP) only

Table 6: Virtual Chassis Summary *(Continued)*

Category	Allowed Routing Engine Members	Allowed Line Card Members
	EX4300 (any models except multigigabit models)	EX4300 (any models except multigigabit models)
	EX4200	EX4200
	EX4500	EX4500
	EX4550	EX4550
	EX3400	EX3400
	EX3300	EX3300
	(Prior to Junos OS Release 18.4R1) EX2300 multigigabit models (EX2300-24MP and EX2300-48MP) only, or EX2300 and EX2300-C (excluding multigigabit models) only	(Prior to Junos OS Release 18.4R1) EX2300 multigigabit models (EX2300-24MP and EX2300-48MP) only, or EX2300 and EX2300-C (excluding multigigabit models) only
	(Starting with Junos OS Release 18.4R1) EX2300, EX2300-C, and EX2300 multigigabit models (EX2300-24MP and EX2300-48MP) (any models in any combination)	(Starting with Junos OS Release 18.4R1) EX2300, EX2300-C, and EX2300 multigigabit models (EX2300-24MP and EX2300-48MP) (any models in any combination))
	EX2200	EX2200

Table 6: Virtual Chassis Summary *(Continued)*

Category	Allowed Routing Engine Members	Allowed Line Card Members
Mixed	QFX5100	QFX5100 QFX3600 QFX3500 EX4300 (any models except multigigabit models)
	QFX3600 QFX3500	QFX3600 QFX3500 EX4300 (any models except multigigabit models)
	EX4600	EX4600 EX4300 (any models except multigigabit models)
	EX4300 multigigabit models (EX4300-48MP)	EX4300 (any models including multigigabit models)
	EX4200, EX4500, EX4550 (any of these switches in any mixed combination)	EX4200, EX4500, EX4550 (any of these switches in any mixed combination)

Understanding the Routing Engine Role in a Virtual Chassis With Different Types of Switches

When you have different types of switches in a Virtual Chassis, the combination of switches you are interconnecting determines which switches should be in the primary Routing Engine role. We also recommend always configuring the same type of switch into the primary and backup Routing Engine roles, to ensure that the switch operating as the primary remains the same type of switch in the event of a switchover.

- In a Virtual Chassis with QFX5110 and QFX5100 switches, which is considered to be a non-mixed QFX5110 Virtual Chassis, we recommend using QFX5110 switches in the primary or backup Routing

Engine roles, and you can use QFX5110 or supported QFX5100 switches for the remaining members in line-card role.

- In a mixed Virtual Chassis with QFX5100 switches combined with QFX3600, QFX3500, or EX4300 switches, you should use QFX5100 switches in the primary Routing Engine role, and you can use QFX5100, QFX3600, QFX3500, or EX4300 switches in the line-card role.
- In a mixed Virtual Chassis with QFX3600 or QFX3500 switches with EX4300 switches, you should use QFX3500 or QFX3600 switches in the Routing Engine role, and you can use QFX3600, QFX3500, or EX4300 switches in the line-card role.
- In a mixed EX4300 and EX4600 Virtual Chassis, EX4600 switches automatically assume the primary and backup Routing Engine roles, and you can use EX4300 or EX4600 switches in the line-card role. (EX4600 switches can be in a mixed Virtual Chassis with any EX4300 model switches except multigigabit models.)
- In a mixed EX4300 Virtual Chassis with EX4300 multigigabit model (EX4300-48MP) switches and any other models of EX4300 switches, you must use EX4300 multigigabit models in the Routing Engine role, and you can use any EX4300 switches in the line-card role.
- In a mixed Virtual Chassis that includes any combination of EX4200 switches, EX4500 switches, and EX4550 switches, any switch can be configured in any role.

In most mixed Virtual Chassis, you must configure your Virtual Chassis to ensure a switch that supports operating as a primary Routing Engine assumes the primary Routing Engine role. Without user configuration, any switch might assume the primary or backup Routing Engine role, with the exception of EX4300 switches in an EX4600 or QFX Series Virtual Chassis, which can never assume the Routing Engine role.

Understanding QFX5100 and QFX5110 Switches in a Virtual Chassis

Up to ten QFX5100 and QFX5110 switches can be interconnected using Virtual Chassis ports (VCPs) to form a *QFX5110 Virtual Chassis*, which is considered to be a non-mixed Virtual Chassis because both types of switches can run the same software image, and you do not need to configure mixed mode. A QFX5110 Virtual Chassis can contain QFX5110 and supported QFX5100 switches in any combination, but we recommend that only QFX5110 switches be in the primary and backup Routing Engine roles, and QFX5100 switches only be configured into line-card role.

NOTE: A QFX5110 Virtual Chassis with QFX5100 switches can only be set up using QFX5110 and QFX5100 switches running the same Junos OS image that includes "-qfx-5e-" in the Junos OS software package filename downloaded from the Software Center. QFX5100 switches running a Junos OS image filename that includes "-qfx-5-" must first be upgraded to the "-qfx-5e-" image to join a QFX5110 Virtual Chassis. (See *Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric*.)

QFX5100 switches can also be interconnected into a mixed Virtual Chassis with QFX3500, QFX3600, and EX4300 switches. See the next section for more information on a mixed QFX5100 Virtual Chassis.

Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Mixed Virtual Chassis

A combination of up to ten EX4300 (excluding multigigabit models), QFX3500, QFX3600, and QFX5100 switches can be interconnected using VCPs to form a mixed QFX Series Virtual Chassis.

In a mixed QFX Series Virtual Chassis that includes QFX5100 switches, only QFX5100 switches should be configured into the Routing Engine role, and the remaining members can be any combination of EX4300, QFX3500, QFX3600, or QFX5100 switches configured into line-card role.

QFX3500 and QFX3600 switches can be interconnected using VCPs to form a Virtual Chassis that is non-mixed; you do not need to configure mixed mode for a Virtual Chassis composed of only QFX3500 and QFX3600 switches. QFX3500 and QFX3600 switches can also be in a mixed QFX Series Virtual Chassis with EX4300 switches. In this combination, only QFX3500 or QFX3600 switches should be configured into the Routing Engine role, and the remaining members can be QFX3500, QFX3600, or EX4300 switches configured into line-card role.

EX4300 switches (excluding multigigabit models) can also be interconnected into a mixed Virtual Chassis with EX4600 switches. See the next section for information on mixed EX4300 and EX4600 Virtual Chassis.

Understanding Mixed EX4300 and EX4600 Virtual Chassis

Any EX4300 switches (except for multigigabit models) and EX4600 switches can be interconnected into a Virtual Chassis.

In a mixed EX4300 and EX4600 Virtual Chassis:

- You can interconnect up to ten member switches.
- An EX4600 switch automatically assumes the primary Routing Engine role.

- EX4300 switches cannot assume the Routing Engine role.

EX4600 switches cannot be in a mixed Virtual Chassis with any other type of switch besides EX4300 switches that are not multigigabit model switches.

Understanding EX4300 Multigigabit and Other EX4300 Model Switches in a Mixed EX4300 Virtual Chassis

You can combine EX4300 multigigabit model (EX4300-48MP) switches with other EX4300 switches into a mixed EX4300 Virtual Chassis.

> **NOTE**: If an EX4300 Virtual Chassis has only EX4300 multigigabit model switches, or only a combination of other models of EX4300 switches, the Virtual Chassis is a non-mixed Virtual Chassis and you should not configure mixed mode.

In a mixed EX4300 Virtual Chassis:

- You can interconnect up to ten member switches.
- You must include the `ieee-clause-82` option when setting mixed mode on the EX4300 switches in the Virtual Chassis that are not multigigabit model switches. This option sets a special port mode (IEEE Clause 82) on the Virtual Chassis ports (VCPs) that enables them to communicate when interconnected with VCPs on EX4300 multigigabit switches. See "Configuring an EX2300, EX3400, or EX4300 Virtual Chassis" on page 80 for details.
- Only EX4300 multigigabit model switches can be in the Routing Engine role.
- EX4300 switches that are not multigigabit model switches must be configured into line-card role.

EX4300 multigigabit switches cannot be in a mixed Virtual Chassis with any other type of switch besides other EX4300 switches.

> **NOTE**: If you remove a non-multigigabit model EX4300 switch from a mixed EX4300 Virtual Chassis with multigigabit model members, you must disable `ieee-clause-82` port mode on the removed switch if you want to reconfigure it as a standalone switch or use it in any other type of mixed Virtual Chassis or any non-mixed Virtual Chassis. Otherwise, the VCPs will not connect with other members in the new Virtual Chassis. See *Removing or Replacing a Member Switch of a Virtual Chassis Configuration* for more details.

Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis

EX4200 switches, EX4500 switches, and EX4550 switches can be interconnected into a Virtual Chassis to form a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, mixed EX4500 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, and EX4550 Virtual Chassis. The mixed Virtual Chassis supports up to 10 member switches regardless of whether the switches are EX4200 switches, EX4500 switches, or EX4550 switches. Any model of EX4200, EX4500, or EX4550 switch can be interconnected into the same mixed Virtual Chassis. The primary-role election process that decides member switch roles in a mixed Virtual Chassis is identical to the primary-role election process in a non-mixed Virtual Chassis, so any member switch in a mixed Virtual Chassis can assume the primary, backup, or line-card role.

EX4200 switches, EX4500 switches, and EX4550 switches cannot be interconnected into a Virtual Chassis with any other switches.

RELATED DOCUMENTATION

Configuring an EX2300, EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis | **80**

Virtual Chassis Overview for Switches | **2**

Understanding EX Series Virtual Chassis | **9**

Understanding QFX Series Virtual Chassis | **25**

Understanding Virtual Chassis Components | **35**

Understanding How the Primary in a Virtual Chassis Is Elected

All switches that are interconnected in a Virtual Chassis configuration are member switches of that Virtual Chassis. Each Virtual Chassis configuration has one member that functions as the *primary* in a Routing Engine role and controls the Virtual Chassis configuration. A second member functions as the backup in the Routing Engine role, and takes control if the primary fails.

When a Virtual Chassis configuration boots, it automatically runs a primary-role election algorithm to determine which member switch assumes the primary role. A Virtual Chassis also applies the same algorithm to choose a new primary or backup member if the member in either role fails.

The first factor the Virtual Chassis considers when choosing the primary member is the *primary-role priority*. On all members, the primary-role priority value is 128 by default. That value can change based on how you provision the Virtual Chassis, as follows:

- In a nonprovisioned Virtual Chassis, you can manually assign primary-role priority values from 0 to 255. A member with primary-role priority 0 will never be elected as primary (or backup), and always stays in the linecard role. In this way, you configure higher primary-role priority values to specify which members can take on the primary (or backup) role.
- In a preprovisioned Virtual Chassis, you can't manually set the primary-role priority on any members. Instead, you assign the Routing Engine role to two member switches. The Virtual Chassis automatically changes the default primary-role priority (128) to 129 on those two members. Then the Virtual Chassis assigns the primary operate in the linecard role by default (you can also explicitly configure them with that role). The Virtual Chassis will never elect a linecard role member as a primary or backup member.

The primary-role election algorithm compares the Virtual Chassis members against the following criteria, in the order listed, until only one member remains under consideration. That member becomes the primary:

1. Choose the member with the highest primary-role priority.
2. Choose the member that was the primary the last time you rebooted the Virtual Chassis.
3. Choose the member that has been in the Virtual Chassis configuration for the longest period of time. (The member switches under consideration must have more than 1 minute between power-up times for this condition to make a difference.)
4. Choose the member with the lowest MAC address.

The primary-role election algorithm does not consider the different switch models or platforms in the Virtual Chassis. For some Virtual Chassis that can contain different types of switches, we require or recommend you configure certain switches in the primary and backup Routing Engine roles. See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details on the types of switches that can be mixed in a Virtual Chassis and which switches can or must be the primary or backup members.

To make sure a specific member is elected as the primary:

1. Power on only the switch that you want to be the primary in the Virtual Chassis.
2. (For a non-provisioned Virtual Chassis) Manually configure the primary-role priorities as follows:
   a. Configure the highest possible primary-role priority value (255) on the member from the first step.

b. On the same member, which is now the primary, configure priority values on the other members. (For example, use the same value or next-highest value on the member you want to be the backup, and lower values on the other members.)

3. (For a preprovisioned Virtual Chassis) Configure the Routing Engine role on the two members that you want to act as the primary and backup members. (You can also explicitly configure the remaining members into the linecard role.)

4. Power on the other members.

You usually want to assign the same (highest) primary-role priority value to the members you want to be the primary and backup Routing Engine members to ensure reliable graceful Routing Engine switchover (GRES) operation. For either non-provisioned or preprovisioned Virtual Chassis, the other election considerations also help keep the primary role from switching back and forth rapidly between the two Routing Engine members under failover conditions.

For more information on configuring a Virtual Chassis, see *Configuring an EX3300 Virtual Chassis (CLI Procedure)*, *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*, "Configuring an EX2300, EX3400, or EX4300 Virtual Chassis" on page 80, or "Configuring an EX4650 or a QFX Series Virtual Chassis" on page 100.

RELATED DOCUMENTATION

Understanding Global Management of a Virtual Chassis

IN THIS SECTION

NOTE: This topic does not apply to EX8200 Virtual Chassis. See *Understanding Global Management of an EX8200 Virtual Chassis*.

A Virtual Chassis is composed of multiple switches, so it has multiple console ports and multiple out-of-band management Ethernet ports located on the switches.

Console Port Session Redirection to Primary Switch

You can connect a PC or laptop directly to a console port of any member switch to set up and configure the Virtual Chassis. When you connect to the console port of any member switch, the console session is redirected to the primary switch, as shown in Figure 1 on page 66.

Figure 1: Console Session Redirection (EX4200 Virtual Chassis Pictured)

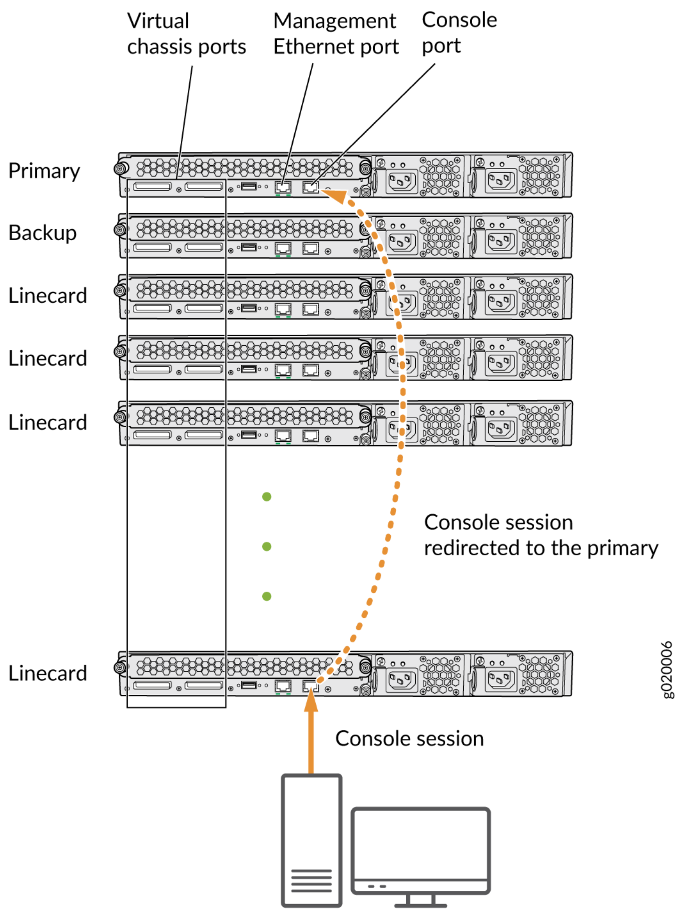

If the primary becomes unavailable, the console session is disconnected from the old primary and a new session is established with the newly elected primary.

An out-of-band management Ethernet port is often referred to simply as a management Ethernet port. It uses a dedicated management channel for device maintenance and allows a system administrator to monitor and manage the switch by remote control.

Logical Port for Virtual Chassis Out-of-Band Management

An out-of-band management Ethernet port is often referred to simply as a management Ethernet port. It uses a dedicated management channel for device maintenance and allows a system administrator to monitor and manage the switch by remote control.

The Virtual Chassis configuration can be managed remotely through SSH or Telnet using a global management interface called the virtual management Ethernet (VME) interface. The VME interface is a *logical interface* representing all of the out-of-band management ports on the member switches. When you connect to the Virtual Chassis configuration using the VME interface's IP address, the connection is redirected to the primary member as shown in Figure 2 on page 67.

Figure 2: Management Ethernet Port Redirection to the VME Interface

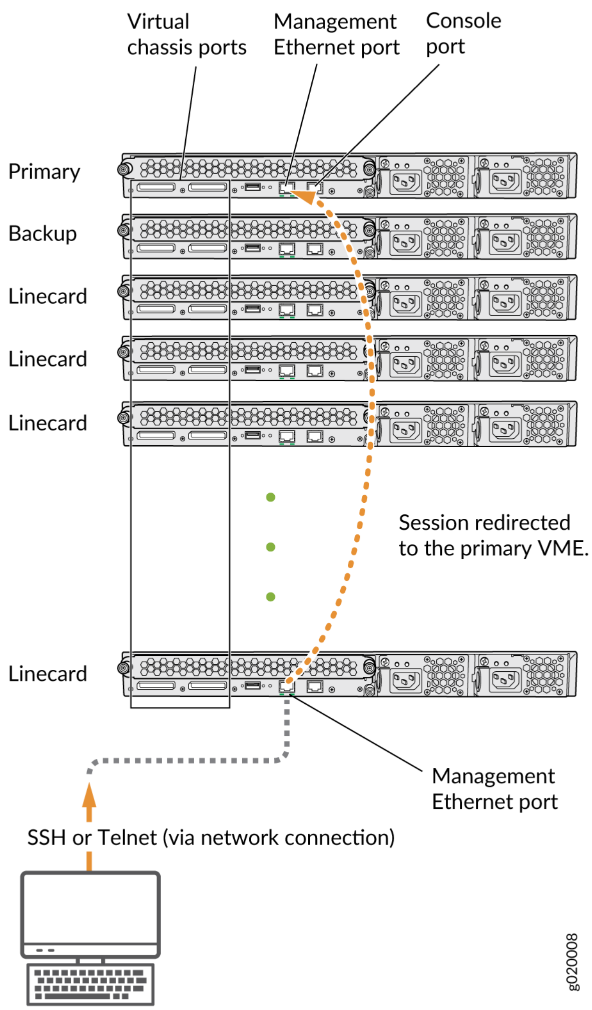

You can configure an IP address for the VME global management interface at any time.

You can perform remote configuration and administration of all members of the Virtual Chassis configuration through the VME interface.

If the primary management Ethernet link is unavailable, the session is redirected through the backup management Ethernet link. If there is no active management Ethernet link on the backup, the VME interface chooses a management Ethernet link on one of the linecard members, selecting the linecard member with the lowest member ID as its first choice.

> **NOTE**: In a QFX Series Virtual Chassis with QFX5110, QFX5120, or QFX5200 switches, the VME interface doesn't redirect properly to the management interface on the backup or linecard member switches if the management link on the primary member switch is down. Instead, you can use the console port on any member switch to globally configure or monitor the Virtual Chassis; this problem doesn't affect console port redirection (see "Console Port Session Redirection to Primary Switch" on page 65).

RELATED DOCUMENTATION

Understanding Virtual Chassis Components | 35

Example: Configuring an EX4200 Virtual Chassis with a Primary and Backup in a Single Wiring Closet

Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)

Understanding Virtual Chassis Port Link Aggregation

IN THIS SECTION

> **NOTE**: This topic applies to all EX Series and QFX Series Virtual Chassis, except for EX8200 Virtual Chassis. See *Understanding Virtual Chassis Port Link Aggregation in an EX8200 Virtual Chassis* for information about EX8200 Virtual Chassis link aggregation.

Within a Virtual Chassis, you can combine physical Ethernet ports to form a logical point-to-point link known as a *link aggregation group (LAG)* or *bundle*. The interfaces that are included in a LAG are sometimes referred to as member interfaces. Do not confuse this term with member switches, which refers to switches that are interconnected into a Virtual Chassis. A LAG in a Virtual Chassis can be composed of member interfaces located in different member switches in the Virtual Chassis, or composed of multiple redundant Virtual Chassis Port (VCP) links between two member switches in the Virtual Chassis, as described in the following sections.

Virtual Chassis Network Interface LAG Among Virtual Chassis Members

When setting up interfaces in a Virtual Chassis, you can configure a combination of physical Ethernet ports belonging to different member switches to form a LAG. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by load-balancing traffic across all available links. If one of the links fails, the system automatically load-balances traffic across all remaining links.

Similarly, if a Virtual Chassis member switch that has LAG member interfaces on multiple member switches fails for any reason, the traffic traversing the LAG can be redirected through the active member switch. This setup has benefits for failover purposes and can be especially beneficial in cases when a member switch needs to be inactive for some time, such as during a software upgrade using NSSU.

NOTE: During an NSSU operation, if you try to view LAG interface status on the primary Routing Engine member using the `show interfaces ae-`*`ae-interface-number`* CLI command, you might see incorrect or zero traffic counts. To work around this problem, run the command on the backup Routing Engine member instead if that member is already loaded and running.

Virtual Chassis Port LAG Between Two Virtual Chassis Members

You can configure optical uplink ports into Virtual Chassis ports (VCPs) that connect EX Series or QFX series switches together to form a Virtual Chassis. When you configure multiple uplink port VCPs connecting the same two member switches, those ports *automatically* form a LAG if the ports are configured to operate at the same link speeds. Each LAG is assigned a positive-integer identifier called a *trunk ID*. Up to 8 redundant VCP links can form a VCP LAG connecting two members in a Virtual Chassis, depending on the number of available ports that can be VCPs.

NOTE: On EX2200 and EX2200-C switches, you can also configure the RJ-45 interfaces, including built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors and 1000BASE-T RJ-45 transceivers, into VCPs. Any interfaces on these switches that are configured as VCPs interconnecting two members will automatically form a LAG, regardless of whether the interfaces are optical transceiver interfaces, RJ-45 transceiver interfaces, or built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors.

A VCP LAG automatically forms when any two member switches are interconnected with two or more VCP links of the same link speed in any of these configurations:

- If the VCP ports on both switches are ports you configured into VCPs or default-configured VCPs (for switches that have default VCPs).
- If the VCP ports on both switches are dedicated VCPs (for switches that have dedicated VCPs).
- In a mixed Virtual Chassis when the VCP links interconnect two different switch models.

A LAG over uplink VCPs provides higher overall bandwidth for forwarding traffic between the member switches connected by the optical VCPs, faster management communications, and greater redundancy of operations among the members than would be available without the LAG.

See "Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port" on page 133 or "Configuring an EX4650 or a QFX Series Virtual Chassis" on page 100 for information about configuring uplink ports into VCPs.

RELATED DOCUMENTATION

Understanding EX Series Virtual Chassis | 9

Understanding Aggregated Ethernet Interfaces and LACP for Switches

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets

Example: Configuring Link Aggregation Groups Using EX4200 Uplink Virtual Chassis Ports

Understanding QFX Series Virtual Chassis | 25

Configuring an EX4650 or a QFX Series Virtual Chassis | 100

Understanding Split and Merge in a Virtual Chassis

IN THIS SECTION

In a Virtual Chassis, you connect two or more switches together to form a unit that is managed as a single chassis. If member switches in the Virtual Chassis fail or you remove member switches, this disrupts the Virtual Chassis configuration. In some situations, the Virtual Chassis configuration splits into two separate Virtual Chassis, which can cause disruptions in the network if the two resulting Virtual Chassis share common resources such as global IP addresses.

The Virtual Chassis split and merge feature is a method to prevent the separate Virtual Chassis configurations from adversely affecting the network. It also enables the two parts to merge back into a single Virtual Chassis configuration.

NOTE: If a Virtual Chassis configuration splits into separate parts, we recommend that you resolve the problem that caused the Virtual Chassis configuration to split as soon as possible.

You can also use this feature to merge two active, separate Virtual Chassis (that have not previously been part of the same configuration) into one Virtual Chassis configuration.

The split and merge feature is enabled by default on EX Series and QFX Series Virtual Chassis. You can disable this feature by using the `set virtual-chassis no-split-detection` command.

What Happens When a Virtual Chassis Configuration Splits

When a Virtual Chassis configuration splits into two separate Virtual Chassis configurations, the individual member switches detect this topology change and run the primary-role election algorithm to select a new primary for each of the two Virtual Chassis configurations. The new primaries then determine whether their Virtual Chassis configuration remains active. One of the configurations remains active based on the following:

- It contains both the stable primary and the stable backup (that is, the primary and backup from the original Virtual Chassis configuration before the split).
- It contains the stable primary and the configuration is greater than half the Virtual Chassis size.
- It contains the stable backup and is at least half the Virtual Chassis size.

In accordance with the rules given in the second and third list items, if the Virtual Chassis configuration splits into two equal parts and the stable primary and stable backup are in different parts, then the part that contains the stable backup becomes active.

> **NOTE**: The number of members in the Virtual Chassis configuration includes all member switches connected to date minus the number whose Virtual Chassis member IDs have been recycled (that is, made available for reassignment). Therefore, the size of the Virtual Chassis configuration increases when a new member switch is detected and decreases when a member switch's ID is recycled.

These rules ensure that only one of the two separate Virtual Chassis configurations created by the split remains active. The member switches in the inactive Virtual Chassis configuration remain in a linecard role. For the inactive members to become active again, one of the following things must happen:

- The problem that caused the original Virtual Chassis configuration to split is resolved, allowing the two Virtual Chassis configurations to merge.
- You load the factory default configuration on the inactive members, which causes the inactive members to function as standalone switches or become part of a different Virtual Chassis configuration.

> **NOTE**: Recycle the member ID if a member (often a linecard or backup) is offline for an extended period of time or removed from the virtual-chassis using the `delete virtual-chassis member <id>` command for pre-provisioned virtual-chasssis and `request virtual-chassis recycle` command for non-provisioned virtual-chassis. This is strongly recommended to prevent any issues which might result in unstable virtual chassis.

Merging Virtual Chassis Configurations

There are two scenarios in which separate Virtual Chassis merge:

- A Virtual Chassis configuration that had split into two is now merging back into a single configuration because the problem that had caused it to split has been resolved.

- You want to merge two Virtual Chassis that had not previously been configured together.

Every Virtual Chassis configuration has a unique ID (VCID) that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a VCID using the `set virtual-chassis id` command. A VCID that you assign takes precedence over automatically assigned VCIDs.

When you reconnect the separate Virtual Chassis configurations or connect them for the first time, the members determine whether or not the separate Virtual Chassis configurations can merge. The members use the following rules to determine whether a merge is possible:

- If the Virtual Chassis configurations have the same VCID, then the configurations can merge. If the two Virtual Chassis were formed as the result of a split, they have the same VCID.
- If the VCIDs are different, then the two configurations can merge only if both are active (inactive configurations cannot merge, ensuring that members removed from one Virtual Chassis configuration do not become members of another Virtual Chassis configuration). If the configurations to merge are both active and one of them has a user-configured VCID, this ID becomes the ID of the merged Virtual Chassis. If neither Virtual Chassis has a user-configured VCID, then the VCID of the configuration with the highest primary-role priority becomes the ID of the merged Virtual Chassis. The resulting merged Virtual Chassis configuration is active.

When you connect two Virtual Chassis configurations:

1. Connecting the two split Virtual Chassis configurations triggers the shortest-path-first (SPF) algorithm. The SPF algorithm computes the network topology and then triggers the primary-role election algorithm. The primary-role election algorithm waits for the members to synchronize the topology information before running.
2. The primary-role election algorithm merges the VCIDs of all the members.
3. Each member runs the primary-role election algorithm to select a primary and a backup from among all members with the same VCIDs. For more information, see "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62.
4. The primary determines whether the Virtual Chassis configuration is active or inactive. (See "What Happens When a Virtual Chassis Configuration Splits" on page 71.)
5. If the Virtual Chassis configuration is active, the primary assigns roles to all members. If the Virtual Chassis configuration is inactive, the primary assigns all members the role of linecard.
6. When the other members receive their role from the primary, they change their role to backup or linecard. They also use the active or inactive state information sent by the primary to set their own state to active or inactive and to construct the Virtual Chassis member list from the information sent by the primary.

7. If the Virtual Chassis state is active, the primary waits for messages from the members indicating that they have changed their roles to the assigned roles, and then the primary changes its own role to primary.

> **NOTE**: When you merge two Virtual Chassis that had not previously been part of the same Virtual Chassis configuration, any configuration settings (such as the settings for Telnet and FTP services, graceful Routing Engine switchover (GRES), fast failover, VLANs, and so on) that exist on the new primary become the configuration settings for all members of the new Virtual Chassis, overwriting any other configuration settings.

RELATED DOCUMENTATION

Disabling Split and Merge in a Virtual Chassis | 140

Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge | 144

Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge

Understanding EX Series Virtual Chassis | 9

Understanding QFX Series Virtual Chassis | 25

Understanding Automatic Software Update on Virtual Chassis Member Switches

IN THIS SECTION

You can use the automatic software update feature to automatically update the Juniper Networks Junos operating system (Junos OS) version on prospective member switches as you add them to an EX Series or QFX Series *Virtual Chassis*.

Automatic Software Update Basics

When you have configured automatic software update on a Virtual Chassis, the Junos OS version is updated on the new member switch when you add it to the Virtual Chassis. The new member switch immediately joins the Virtual Chassis configuration and is put in the active state.

For a standalone switch to join an existing Virtual Chassis, it must be running the same version of Junos OS that is running on the Virtual Chassis primary. When the primary in a Virtual Chassis detects that a new switch has been added to the configuration, it checks the software version on the new switch. If the software version on the new switch is not the same as the version running on the primary, the primary keeps the new switch in the inactive state. If you have not enabled the automatic software update feature, you have to manually install the correct software version on each prospective member switch as it is added to the Virtual Chassis.

Automatic Software Update Restrictions

Refer to Feature Explorer to see the EX Series and QFX Series Virtual Chassis that support automatic software updates, and the Junos OS release versions where the feature was introduced on each platform. Where the automatic software update feature is supported, see the Junos OS Release Notes for the release version running on the primary for any limitations in using automatic software updates between that release version and any prospective member switch release versions.

> **CAUTION**: A QFX5100 switch running a Junos OS software image with "-qfx-5-" in the package filename *must* first be upgraded to a Junos OS software image with "-qfx-5e-" in the package filename before it can be added to a QFX5110 Virtual Chassis or VCF. The automatic software update process cannot update a switch from a "-qfx-5-" image to a "-qfx-5e-" image. See *Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric*.
>
> After a QFX5100 switch is installed with a "-qfx-5e-" Junos OS software image, the automatic software update process can successfully update the switch automatically with a different version of a "-qfx-5e-" Junos OS image to match the other members in the Virtual Chassis or VCF.

RELATED DOCUMENTATION

Understanding Software Upgrades in a Virtual Chassis | 170

Configuring Automatic Software Update on Virtual Chassis Member Switches | 141

Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches

Understanding MAC Address Assignment on a Virtual Chassis

In a *Virtual Chassis*, multiple switches—each with its own set of interfaces with unique MAC addresses—are connected together to form one chassis that can be managed as a single switch. The MAC address assigned to each network-facing interface on the switch changes when the switch joins a Virtual Chassis. Because all Layer 2 traffic decisions are based on an interface's MAC address, understanding MAC address assignment is important to understanding how network traffic is forwarded and received by the Virtual Chassis. For additional information about how a network uses MAC addresses to forward and receive traffic, see Understanding Bridging and VLANs on Switches.

When a Virtual Chassis is formed, the MAC address of the switch in the primary role becomes the system MAC base address. The Virtual Chassis assigns the system MAC base address as the MAC address for all Layer 3 interfaces within the Virtual Chassis. The Virtual Chassis also assigns the system MAC base address to the virtual management Ethernet (VME) interface and to all of the virtual LANs (VLANs) in the Virtual Chassis.

The system MAC base address does not change in the event of a switchover if the switch that was originally configured in the primary role remains a member of the Virtual Chassis. If the switch that was originally configured in the primary role is removed from the Virtual Chassis, the MAC address of the current member switch in the primary role is assigned as the system MAC base address after the MAC persistence timer interval has expired. You can configure the MAC persistence timer interval.

For Layer 2 and aggregated Ethernet interfaces, the Virtual Chassis assigns a unique MAC address that is derived from the member switch MAC address to each interface. The assignment of a unique MAC address to each network interface helps ensure that functions that require MAC address differentiation—such as redundant trunk groups (RTGs), Link Aggregation Control Protocol (LACP), and general monitoring functions—can function properly.

NOTE: Unique MAC address assignment for Layer 2 and aggregated Ethernet interfaces in a Virtual Chassis was introduced in Junos OS Release 11.3. The same MAC address could be assigned to interfaces on different member switches in the same Virtual Chassis prior to this release.

If you reconfigure a Layer 2 interface into a Layer 3 interface, or the reverse, within a Virtual Chassis, the MAC address of that interface changes accordingly.

MAC addresses are assigned to interfaces in a Virtual Chassis automatically—no user configuration is possible or required. You can view the MAC addresses that are assigned to the interfaces by using the `show interfaces` command.

RELATED DOCUMENTATION

Understanding MAC Address Assignment on an EX Series Switch

Configuring the Timer for the Backup Member to Start Using Its Own MAC Address as Primary of a Virtual Chassis | **132**

Understanding EX Series Virtual Chassis | **9**

EX8200 Virtual Chassis Overview

Understanding QFX Series Virtual Chassis | **25**

Understanding High Availability on an EX Series Virtual Chassis

You increase your network's high availability (HA) when you interconnect a Juniper Networks EX Series Ethernet switch into a *Virtual Chassis*. A Virtual Chassis is more fault tolerant then a standalone EX series switch because it remains up when a single member switch fails, and provides sub-second convergence in the case of a device or link failure.

You can further improve HA by configuring the HA features available for your EX Series Virtual Chassis. You can, for instance, configure Link Aggregation Groups (LAG) bundles to include member links on multiple member switches in the same Virtual Chassis. This configuration increases fault tolerance because traffic traversing the LAG can be redirected to an active member switch when a single member switch fails.

A Virtual Chassis has dual Routing Engines—the switch in the primary role and the switch in the backup role—and therefore supports many HA features not supported on standalone EX Series switches, such as Graceful Routing Engine Switchover (GRES) for hitless failover. For information on which of the High Availability features listed in Table 7 on page 78 are supported in your EX Series Virtual Chassis, see Feature Explorer.

Many HA features for the EX Series Virtual Chassis are designed to improve network resiliency after a Routing Engine switchover. Table 7 on page 78 describes the effects of a Routing Engine switchover when no high availability features are enabled and when some High Availability features are enabled.

Table 7: Effects of a Routing Engine Switchover

High Availability Feature	Effect of Routing Engine Switchover
No HA features enabled	Kernel and forwarding state information is not preserved to the backup Routing Engine. A convergence process that requires all interfaces on the Virtual Chassis to be taken offline has to be performed before the Virtual Chassis returns online. The switchover can take several minutes and the Virtual Chassis does not send or receive traffic until the switchover is complete.
Graceful Routing Engine switchover (GRES) enabled	Kernel and forwarding state information is preserved on both Routing Engines, so the convergence process does not occur and the switchover happens quickly with minimal traffic loss.
Nonstop active routing (NSR), Nonstop bridging (NSB), or both enabled	Layer 2 protocols that are supported by NSB are not disrupted by a Routing Engine switchover when NSB is enabled. Layer 2 protocol information for all active Layer 2 protocols is stored on both Routing Engines when NSB is enabled. Layer 3 protocols that are supported by NSR are not disrupted by a Routing Engine switchover when NSR is enabled. Layer 3 protocol information for all active Layer protocols is stored on both Routing Engines when NSR is enabled.
Graceful Protocol Restart enabled	Traffic is not interrupted during the switchover. Interface and kernel information is preserved. Graceful restart protocol extensions quickly collect and restore routing information for supported protocols from the neighboring devices.

RELATED DOCUMENTATION

Understanding EX Series Virtual Chassis | 9

2

CHAPTER

Virtual Chassis Configuration

Configuring an EX2300, EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis

IN THIS SECTION

You can use the procedures in this topic to configure:

- An EX2300 Virtual Chassis

> **NOTE**: Junos OS releases prior to 18.4R1 support forming an EX2300 Virtual Chassis using only EX2300 multigigabit switches or only EX2300 switches that are not multigigabit model switches.
>
> Starting in Junos OS Release 18.4R1, EX2300, EX2300-C, and EX2300 multigigabit switches can all be combined in the same non-mixed Virtual Chassis.

- An EX3400 Virtual Chassis
- A non-mixed EX4300 Virtual Chassis

> **NOTE**: An EX4300 Virtual Chassis operates as a non-mixed Virtual Chassis if it is composed of only EX4300 multigigabit model switches, or composed only of any combination of any other EX4300 switches excluding the multigigabit models.
>
> You must configure mixed mode if you combine EX4300 multigigabit models with other EX4300 models in an EX4300 Virtual Chassis.

- A mixed EX4300 Virtual Chassis with EX4300 multigigabit model (EX4300-48MP) switches interconnected with other EX4300 model switches
- An EX4100 Virtual Chassis

> **NOTE**: You can combine any models of EX4100 switches, including EX4100 multigigabit models, and EX4100-F into a Virtual Chassis without needing to configure mixed mode.

- An EX4400 Virtual Chassis

> **NOTE**: You can combine any models of EX4400 switches, including multigigabit models, into an EX4400 Virtual Chassis without needing to configure mixed mode.

You can mix non-multigigabit model EX4300 switches with other switches in a Virtual Chassis or Virtual Chassis Fabric (VCF) in the following supported combinations. In these cases, use the following configuration procedures instead of the procedures in this topic:

- A mixed EX4600 Virtual Chassis that includes EX4600 and EX4300 member switches: "Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis" on page 93.
- A mixed QFX Series Virtual Chassis that includes supported QFX Series and EX4300 switches: "Configuring an EX4650 or a QFX Series Virtual Chassis" on page 100.
- A mixed VCF with EX4300 switches as leaf nodes: *Preprovisioning a Virtual Chassis Fabric* or *Autoprovisioning a Virtual Chassis Fabric*.

Use the following requirements and guidelines to plan the devices to include in your Virtual Chassis:

- EX2300 switches:

  In Junos OS releases prior to 18.4R1, you can interconnect EX2300 and EX2300-C switches into a Virtual Chassis, or you can interconnect EX2300 multigigabit model switches (EX2300-24MP and EX2300-48MP) into a Virtual Chassis. You can’t combine EX2300 or EX2300-C switches with EX2300 multigigabit model switches in a Virtual Chassis.

  Starting in Junos OS Release 18.4R1, you can combine EX2300, EX2300-C, and EX2300 multigigabit switches in the same non-mixed Virtual Chassis, and use any of these switches in any role (primary Routing Engine role, backup Routing Engine role, or linecard role).

- EX3400 switches:

  You can interconnect EX3400 switches only with other EX3400 switches in a Virtual Chassis (no mixed mode).

- EX4300 switches:

  You can interconnect EX4300 switches excluding multigigabit models into a *non-mixed* EX4300 Virtual Chassis.

You can also connect EX4300 multigigabit model switches (EX4300-48MP) together into a *non-mixed* EX4300 Virtual Chassis.

You can combine EX4300 multigigabit model switches with other EX4300 model switches as a *mixed* EX4300 Virtual Chassis with the following configuration:

- You must configure the Virtual Chassis into mixed mode.
- You must also include a special port mode option (`ieee-clause-82`) when you configure mixed mode on the EX4300 switches that are not multigigabit models. This port mode enables the Virtual Chassis ports (VCPs) on EX4300 non-multigigabit model switches to communicate with VCPs on multigigabit model members.

  NOTE: If you remove a non-multigigabit model EX4300 switch from a mixed EX4300 Virtual Chassis with multigigabit model members, remember to disable `ieee-clause-82` port mode on the removed switch if you want to reconfigure it as a standalone switch or use it in any other type of mixed Virtual Chassis or non-mixed Virtual Chassis. Otherwise, the VCPs will not connect with other members in the new Virtual Chassis. (See *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*.)

- The members in the Routing Engine role must be multigigabit model (EX4300-48MP) switches.

- EX4100 switches:

  You can interconnect EX4100 switches only with other EX4100 switches in a Virtual Chassis, including EX4100 multigigabit models, and EX4100-F model. You do not need to configure mixed mode when combining EX4100 multigigabit models with other EX4100 and EX4100-F models.

- EX4400 switches:

  You can interconnect EX4400 switches only with other EX4400 switches in a Virtual Chassis, including EX4400 multigigabit models. You do not need to configure mixed mode when combining EX4400 multigigabit models with other EX4400 models.

Use these guidelines to plan the VCP connections:

- In a non-mixed EX4300 Virtual Chassis with only EX4300 multigigabit model (EX4300-48MP) switches, interconnect the member switches using the dedicated VCPs—the 40-Gpbs QSFP+ ports on the rear panel. These are the only ports that can be used as VCPs on EX4300-48MP switches.
- In a mixed EX4300 Virtual Chassis with a combination of EX4300 multigigabit model switches and other EX4300 model switches, you must use 40-Gbps QSFP+ ports on the other EX4300 model switches as VCPs and interconnect those ports with the dedicated VCPs on the multigigabit model switch members.

NOTE: All QSFP+ ports on EX4300 non-multigigabit switches are configured as VCPs in the default factory configuration.

- In non-mixed EX2300, EX3400, and non-multigigabit model EX4300 Virtual Chassis, use uplink ports that you configure as VCPs or that are VCPs by default to interconnect the member switches. Keep the following in mind about VCPs on these switches:

  - EX2300 switches do not have any ports that are configured by default as VCPs. You must explicitly configure the ports you want to use as VCPs.

  - The QSFP+ uplink ports on EX3400 and EX4300 switches support 40-Gbps speeds. These ports are set as VCPs by default, so you don't need to explicitly configure them.

  - You can configure the SFP+ uplink ports on any of these switches as VCPs. These ports support 10-Gbps speeds and can connect switches that are up to 6.2 miles (10 km) apart.

    NOTE: The only exceptions are the four *built-in* 10-Gbps SFP+ ports on 32-port EX4300 switches. You can't use the built-in ports as VCPs.

    Also, for uplink ports on these switches that support SFP or SFP+ transceivers, you can't form a Virtual Chassis using ports that have SFP transceivers installed. The ports must have SFP+ transceivers installed for them to function properly as VCPs.

- The simplest way to interconnect EX3400 or EX4300 switches into a non-mixed EX3400 or EX4300 Virtual Chassis is to interconnect them into a Virtual Chassis by using the QSFP+ ports (the default VCPs).

  For an EX3400 or EX4300 Virtual Chassis, if you are using the QSFP+ ports for another purpose, or for any EX2300 Virtual Chassis, you must configure SFP+ uplink module ports into VCPs.

- In an EX4400 Virtual Chassis, you must use the default VCPs to interconnect the member switches, which are the only ports that can be used as VCPs on any model of these switches. The default VCPs on the EX4400 and multigigabyte models are the two 100-Gbps ports on the rear panel, which operate as two logical 50-Gbps VCPs each for a total of four logical VCP interfaces on the switch. If you previously converted the default VCPs into network ports, you must convert them back into VCPs using the request virtual- chassis mode network-port disable command. You must then reboot the switch for the port mode conversion to take effect.

- In an EX4100 Virtual Chassis, you must use the default VCPs to interconnect the member switches, which are the only ports that can be used as VCPs on any model of these switches. The default VCPs on the EX4100 and EX4100 multigigabyte models are the four 25-Gbps ports on the front panel. If you previously converted the default VCPs into network ports, you must convert them back into

VCPs using the request virtual- chassis mode network-port disable command. You must then reboot the switch for the port mode conversion to take effect.

- In an EX4100-F Virtual Chassis, you must use the default VCPs to interconnect the member switches, which are the only ports that can be used as VCPs on any model of these switches. The default VCPs on the EX4100-F models are the four 10-Gbps ports on the front panel. If you previously converted the default VCPs into network ports, you must convert them back into VCPs using the request virtual- chassis mode network-port disable command. You must then reboot the switch for the port mode conversion to take effect.

- If you need additional VCP bandwidth between two member switches, you can configure additional ports as VCPs and create redundant links between the member switches.

  Redundant VCP links are not required to be the same speed, but the links with identical speeds automatically form a VCP link aggregation group (LAG) to provide resiliency to the Virtual Chassis. For example, if you have two 40-Gbps QSFP+ ports and two 10-Gbps SFP+ ports configured as VCPs connecting the same two member switches to each other, the member switches form two LAGs—one LAG with two 40-Gbps QSFP+ port links and another with two 10-Gbps SFP+ port links.

NOTE: A Virtual Chassis configuration has two Routing Engines—the primary switch and the backup switch. We recommend that you always use `commit synchronize` rather than simply `commit` to save configuration changes made for a Virtual Chassis. This ensures that you save the configuration changes on both Routing Engines at the same time.

You can configure the Virtual Chassis using either of the following options:

- A nonprovisioned configuration—The primary sequentially assigns a member ID to other member switches. The role is determined by the primary-role priority value and other factors in the primary-role election algorithm.

- A preprovisioned configuration—You can deterministically control the member ID and role assigned to a member switch by tying the member switch to its serial number.

All member switches must be running the same version of Junos OS to form a Virtual Chassis.

NOTE: You must configure a VLAN on all interfaces in the Virtual Chassis, except on member switch 0, before the interfaces can send or receive traffic. This is because the interfaces on member switch 0 are initially placed into the default VLAN, but the interfaces on all other member switches are not placed into any VLAN. See Configuring VLANs for EX Series Switches with ELS Support (CLI Procedure).

NOTE: On an EX4300 Virtual Chassis, STP is disabled on all interfaces except the interfaces on member switch 0 until some type of spanning-tree protocol is enabled. See Configuring RSTP on EX Series Switches (CLI Procedure) (RSTP is the default spanning-tree protocol), Configuring MSTP on Switches, or Configuring VSTP Protocol to enable a spanning-tree protocol on the interfaces in your EX4300 Virtual Chassis.

Configuring an EX2300, EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis with a Nonprovisioned Configuration File

You can use a nonprovisioned configuration to configure an EX2300, EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis.

This procedure shows example configuration steps for a Virtual Chassis with two to ten members. You can have up to 4 members in an EX2300 Virtual Chassis, and up to 10 members in an EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis.

NOTE: We recommend that you physically cable the interconnecting ports as the final step of this procedure.
You can, however, configure the Virtual Chassis while the cables are physically connected.

1. Power on only the switch that you will use as the primary switch.

   NOTE: For a mixed EX4300 Virtual Chassis with EX4300 multigigabit model (EX4300-48MP) switches and other EX4300 model switches, the members in the primary and backup Routing Engine roles must be EX4300 multigigabit model switches.

2. (Required for a mixed EX4300 Virtual Chassis only) Set the primary switch into mixed mode, and reboot the switch for the change to take effect:

   ```
   user@device> request virtual-chassis mode mixed reboot
   ```

3. If you are configuring a mixed Virtual Chassis, wait for the reboot to complete before performing this step. Run the EZSetup program on the primary switch, specifying the identification parameters. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* for details.

NOTE: The properties that you specify for the primary switch apply to the entire Virtual Chassis configuration.

4. (Optional) Configure the primary switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis (see "Understanding Global Management of a Virtual Chassis" on page 64):

```
[edit]
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```

5. (Optional) Configure primary-role priority for the other member switches. For example, for a ten-member Virtual Chassis:

```
[edit virtual-chassis]
user@switch# set member 0 mastership-priority 255
user@switch# set member 1 mastership-priority 255
user@switch# set member 2 mastership-priority 10
user@switch# set member 3 mastership-priority 9
user@switch# set member 4 mastership-priority 8
user@switch# set member 5 mastership-priority 7
user@switch# set member 6 mastership-priority 6
user@switch# set member 7 mastership-priority 5
user@switch# set member 8 mastership-priority 4
user@switch# set member 9 mastership-priority 3
```

The primary-role priority value determines the roles in a non-provisioned Virtual Chassis configuration. The switches with the highest primary-role priority values assume the primary and backup roles. All other switches assume the linecard role.

If you do not configure the primary-role priority for any switch in your Virtual Chassis, including when you do not configure the Virtual Chassis, all switches assume the default primary-role priority of 128. The primary-role election algorithm selects the roles for the member switches. In most cases, the switches that have been powered on the longest assume the primary and backup roles when all Virtual Chassis member switches are configured with the same primary-role priority. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62 for additional information on the primary-role election algorithm.

A switch with a primary-role priority of 0 never assumes the primary or backup role.

NOTE: We recommend that you specify the same primary-role priority value for the intended primary and backup members.

6. (Optional: Recommended for a two-member Virtual Chassis) On the primary switch, disable the split and merge feature:

```
[edit virtual-chassis]
user@switch# set no-split-detection
```

7. Commit the configured items.
8. Power on the other member switches.
9. (Required for a mixed EX4300 Virtual Chassis only) Set each additional individual EX4300 multigigabit model (EX4300-48MP) switch into mixed mode and reboot the switch for the change to take effect:

```
user@device> request virtual-chassis mode mixed reboot
```

   Set each of the other EX4300 switches that are not multigigabit model switches into mixed mode with the `ieee-clause-82` option, and reboot the reboot the switch for the change to take effect:

```
user@device> request virtual-chassis mode ieee-clause-82 mixed reboot
```

10. If needed, on each individual member switch, configure SFP+ uplink ports that will be used as VCPs to interconnect the member switches.

    This step is not needed for:

    - EX4300 multigigabit model (EX4300-48MP) switches, which have dedicated VCPs on the rear panel of the switch. Dedicated VCPs do not require configuration, and you must use the dedicated ports on these switches because no other ports on these switches are supported as VCPs.

    - EX3400 switches and EX4300 member switches that are not multigigabit model switches, if you are using the QSFP+ ports that are VCPs by default. You only need to configure a QSFP+ port as a VCP if you previously configured the QSFP+ port into a network port. In that case, you can perform this step to configure the QSFP+ port back into a VCP.
11. (EX4100, EX-4100-F, and EX4400 switches only) If the ports that you will use as VCPs were previously converted into network ports on any member switches, convert them back into VCPs in this step using the `request virtual-chassis mode network-port disable` command, and reboot the affected switches for the change to take effect. The `request virtual-chassis vc-port` command shown in Step 10 doesn't convert a network port into a VCP on EX4100, EX4100-F, and EX4400 switches.

> **NOTE**: This command changes the port mode for all of the VCPs on the switch. The four ports must both operate together as VCPs or as network ports. You are also required to reboot the switch for any mode command changes to take affect. You can optionally include the reboot option with the mode command to reboot the switch immediately, as shown here. Otherwise, you can alternatively reboot the switch later with a separate reboot command.

For example:

```
user@switch> request virtual-chassis mode network-port disable reboot
```

12. Cable the ports interconnecting the members if they were not connected earlier.

> **NOTE**: If you want to change the member ID that the primary has assigned to a member switch, use the `request virtual-chassis renumber` command.

Configuring an EX2300, EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis with a Preprovisioned Configuration File

Preprovisioning a Virtual Chassis configuration allows you to assign the member ID and role for each switch in the Virtual Chassis.

This procedure shows example configuration steps applicable to a Virtual Chassis with two to ten members; up to 4 members are supported in an EX2300 Virtual Chassis, and up to 10 members are supported in an EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis.

To configure a Virtual Chassis using a preprovisioned configuration:

> **NOTE**: We recommend that you physically cable the optical ports as the final step of this procedure.
> You can, however, configure the Virtual Chassis while the cables are physically connected.

1. Make a list of the serial numbers of all the switches to be connected in a Virtual Chassis configuration.

> **NOTE**: Serial number values are case-sensitive.

2. Note the intended role (`routing-engine` or `line-card`) of each switch. If you configure the member with a `routing-engine` role, it is eligible to function in the primary or backup role. If you configure the member with a `line-card` role, it is not eligible to function in the primary or backup role.

> **NOTE**: For a mixed EX4300 Virtual Chassis with EX4300 multigigabit model (EX4300-48MP) switches and other EX4300 model switches, the members in the primary and backup Routing Engine roles must be EX4300 multigigabit model switches.

3. Power on only the switch that you plan to use as the primary switch.
4. (Required for a mixed EX4300 Virtual Chassis only) Set the primary switch into mixed mode, and reboot the switch for the change to take effect:

```
user@device> request virtual-chassis mode mixed reboot
```

5. If you are configuring a mixed Virtual Chassis, wait for the reboot to complete before performing this step. Run the EZSetup program on the primary switch, specifying the identification parameters. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* for details.

> **NOTE**: The properties that you specify for the primary switch apply to the entire Virtual Chassis configuration.

6. (Optional) Configure the primary switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis (see "Understanding Global Management of a Virtual Chassis" on page 64):

```
[edit]
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```

7. Specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@switch# set preprovisioned
```

8. Specify all the members that you want included in the Virtual Chassis, listing each switch's serial number with the desired member ID and role. For example, for a ten-member Virtual Chassis:

NOTE: You can retrieve the switch's serial number using the **show chassis hardware** command output or by viewing the serial number ID label on the switch. See *Locating the Serial Number on an EX2300 Switch or Component*, *Locating the Serial Number on an EX3400 Switch or Component*, or *Locating the Serial Number on an EX4300 Switch or Component*. Serial number values are case-sensitive.

```
[edit virtual-chassis]
user@switch# set member 0 serial-number abc123 role routing-engine
user@switch# set member 1 serial-number def456 role routing-engine
user@switch# set member 2 serial-number ghi789 role line-card
user@switch# set member 3 serial-number jkl012 role line-card
user@switch# set member 4 serial-number mno013 role line-card
user@switch# set member 5 serial-number pqr014 role line-card
user@switch# set member 6 serial-number stu015 role line-card
user@switch# set member 7 serial-number vwx016 role line-card
user@switch# set member 8 serial-number yzz017 role line-card
user@switch# set member 9 serial-number aaa018 role line-card
```

9. (Optional: Recommended for a two-member Virtual Chassis) Disable the split and merge feature:

```
[edit virtual-chassis]
user@switch# set no-split-detection
```

10. Commit the configured items.
11. Power on the other member switches. The member IDs and roles have been determined by the configuration, so you can power on the member switches in any order.
12. (Required for a mixed EX4300 Virtual Chassis only) Set each additional individual EX4300 multigigabit model (EX4300-48MP) switch into mixed mode and reboot the switch for the change to take effect:

```
user@device> request virtual-chassis mode mixed reboot
```

Set each of the other EX4300 switches that are not multigigabit model switches into mixed mode with the `ieee-clause-82` option, and reboot the reboot the switch for the change to take effect:

```
user@device> request virtual-chassis mode ieee-clause-82 mixed reboot
```

13. If needed, on each individual member switch, configure the SFP+ uplink ports that will be used as VCPs to interconnect the member switches.

    This step is not needed for:

    - EX4300 multigigabit model (EX4300-48MP) switches, which have dedicated VCPs on the rear panel of the switch. Dedicated VCPs do not require configuration, and you must use the dedicated ports on these switches because no other ports on these switches are supported as VCPs.

    - EX3400 switches and EX4300 member switches that are not multigigabit model switches, if you are using the QSFP+ ports that are VCPs by default. You only need to configure a QSFP+ port as a VCP if you previously configured the QSFP+ port into a network port. In that case, you can perform this step to configure the QSFP+ port back into a VCP.

    - EX4400 switches, which have the rear panel 100 Gigabit Ethernet QSFP28 ports set as VCPs by default. Those are the only ports that you can use as VCPs on these switches. If you previously converted the default VCPs into network ports, you can't convert them back into VCPs using the command in this step. See Step 15 instead.

    Use the *request virtual-chassis vc-port* command to configure SFP+ or QSFP+ ports into VCPs. For example, for a four-member Virtual Chassis:

```
user@switch-0> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-0> request virtual-chassis vc-port set pic-slot 1 port 1
```

```
user@switch-1> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-1> request virtual-chassis vc-port set pic-slot 1 port 1
```

```
user@switch-2> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-2> request virtual-chassis vc-port set pic-slot 1 port 1
```

```
user@switch-3> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-3> request virtual-chassis vc-port set pic-slot 1 port 1
```

14. (EX4100, EX-4100-F switches only) If the ports that you will use as VCPs were previously converted into network ports on any member switches, convert them back into VCPs in this step using the `request virtual-chassis mode network-port disable` command, and reboot the affected switches for the change to take effect. The `request virtual-chassis vc-port` command shown in Step 10 doesn't convert a network port into a VCP on EX4100 and EX4100-F switches.

 > **NOTE**: This command changes the port mode for all of the VCPs on the switch. The four ports must both operate together as VCPs or as network ports. You are also required to reboot the switch for any mode command changes to take affect. You can optionally include the reboot option with the mode command to reboot the switch immediately, as shown here. Otherwise, you can alternatively reboot the switch later with a separate reboot command.

 For example:

 ```
 user@switch> request virtual-chassis mode network-port disable reboot
 ```

15. EX4400 switches only) If the ports that you will use as VCPs were previously converted into network ports on any member switches, convert them back into VCPs in this step using the `request virtual-chassis mode network-port disable` command, and reboot the affected switches for the change to take effect. The `request virtual-chassis vc-port` command shown in Step 13 doesn't convert a network port into a VCP on EX4400 switches.

 EX4400 switches have the two rear panel 100 Gigabit Ethernet QSFP28 ports configured into four logical 50-Gbps VCP interfaces by default. These are the only ports that you can use to connect member switches into an EX4400 Virtual Chassis. If you previously converted them to network port mode, you must disable network port mode to restore them to the default VCP mode for the Virtual Chassis to form when you cable the member switches together. You can use the `show virtual-chassis mode` command to check whether the switch has network port mode enabled or not.

 > **NOTE**: This command changes the port mode for all of the VCPs on the switch. The two ports must both be set together to VCP port mode or to network port mode. You are also required to reboot the switch for any mode command changes to take affect. You can optionally include the reboot option with the mode command to reboot the switch immediately, as shown here. (Otherwise, you can alternatively reboot the switch later with a separate reboot command.)

For example:

```
user@switch> request virtual-chassis mode network-port disable reboot
```

16. Cable the ports interconnecting the members if they were not connected earlier.

NOTE: You cannot modify the primary-role priority when you are using a preprovisioned configuration. The primary-role priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two Routing Engines are assigned the same primary-role priority value. However, the member that was powered on first has higher prioritization according to the primary-role election algorithm. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62.

RELATED DOCUMENTATION

Configuring Primary Role of a Virtual Chassis | 129

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

IN THIS SECTION

This topic provides information about configuring a non-mixed EX4600 Virtual Chassis or a mixed EX4600 Virtual Chassis that includes EX4600 switches and EX4300 switches. It does not cover EX4650 Virtual Chassis, which is more like a QFX5120 Virtual Chassis than an QFX4600 Virtual Chassis; instead, see "Configuring an EX4650 or a QFX Series Virtual Chassis" on page 100.

NOTE: A mixed EX4600 Virtual Chassis can include any models of EX4300 switches except multigigabit EX4300 models (EX4300-48MP).

You configure 40-Gbps QSFP+ and 10-Gbps SFP+ uplink ports as Virtual Chassis ports (VCPs) to interconnect members in an EX4600 Virtual Chassis. Uplink ports can connect switches that are several miles apart in different buildings into the same Virtual Chassis.

You must configure QSFP+ or SFP+ uplink module ports into VCPs to create a non-mixed or mixed EX4600 Virtual Chassis. If you need additional VCP bandwidth between two member switches, you can configure additional ports as VCPs between the member switches. The ports that have identical speeds become links in a link aggregation group (LAG) to provide resiliency to the Virtual Chassis; for instance, if you had two QSFP+ ports and two SFP+ ports configured as VCPs connecting to the same switch, you would have two LAGs—one LAG with two 40Gbps QSFP+ port member links and another with two 10Gbps SFP+ port member links—between the member switches.

NOTE: A Virtual Chassis configuration has two Routing Engines—the primary switch and the backup switch. We recommend that you always commit Virtual Chassis changes by using the `commit synchronize` command rather than `commit`. This ensures that the configuration changes are saved on both Routing Engines at the same time.

You can set up a EX4600 Virtual Chassis with either:

- A nonprovisioned configuration—The primary sequentially assigns a member ID to other member switches, and determines the role from the primary-role priority value and other factors in the primary-role election algorithm.
- A preprovisioned configuration—You deterministically control the member ID and role assigned to a member switch by tying the member switch to its serial number.

Configuring an EX4600 Virtual Chassis with a Nonprovisioned Configuration File

You can use nonprovisioned configuration to configure an EX4600 Virtual Chassis.

To configure the Virtual Chassis using a nonprovisioned configuration:

NOTE: You can configure a EX4600 Virtual Chassis while the cables are or are not physically connected.

1. Power on only the switch that you plan to use as the primary switch.

   If you are configuring a mixed EX4300 and EX4600 Virtual Chassis, you must use an EX4600 switch. We don't support a mixed EX4300 and EX4600 Virtual Chassis with an EX4300 in the primary routing engine role.
2. (Required for a mixed Virtual Chassis only) Set the primary switch into mixed mode, and reboot the switch for the change to take effect:

```
user@device> request virtual-chassis mode mixed reboot
```

3. If you are configuring a mixed Virtual Chassis, wait for the reboot to complete before performing this step.Specify the identification parameters for the switch by completing the initial configuration. See *Configuring Junos OS on the EX4600*.

   NOTE: The properties that you specify for the primary switch apply to the entire Virtual Chassis configuration.

4. (Optional) Configure the primary switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis:

```
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```

5. (Required for mixed EX4300 and EX4600 Virtual Chassis, optional for other Virtual Chassis) Configure primary-role priority for the member switches:

   NOTE: If you are configuring a mixed EX4300 and EX4600 Virtual Chassis, configure the EX4600 switches with the highest primary role priorities to ensure EX4300 switches do not assume the primary role. We don't support a mixed EX4300 and EX4600 switch operating with an EX4300 switch in the primary role.

```
[edit virtual-chassis]
user@switch# set member 0 mastership-priority 255
user@switch# set member 1 mastership-priority 255
```

6. (Optional. Recommended for a two-member Virtual Chassis) On the primary switch, disable the split and merge feature:

```
[edit virtual-chassis]
user@switch# set no-split-detection
```

7. Power on the other member switches.
8. (Required for a mixed Virtual Chassis only) Set each individual switch into mixed mode, and reboot the switch for the change to take effect:

```
user@device> request virtual-chassis mode mixed reboot
```

9. On each individual member switch, use the following command to configure the ports that will be used to interconnect the member switches into VCPs:

```
user@switch> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
local
```

where *pic-slot-number* is the PIC slot number.

For instance, if you wanted to set port 0 on the QSFP+ interface on PIC slot 2 as a VCP:

```
user@switch> request virtual-chassis vc-port set pic-slot 2 port 0 local
```

The VCPs automatically bundle into a Link Aggregation Group when two or more interfaces of the same speed are configured into VCPs between the same two member switches. See "Understanding Virtual Chassis Port Link Aggregation" on page 68.

NOTE: By default, a Virtual Chassis forms with a nonprovisioned configuration if you don't set the member roles and primary role priorities. In a non-provisioned configuration, the primary-role priority value for each member switch is `128`, and the primary role is selected by default. You can modify the primary-role priority to change the role a member takes on. See "Configuring Primary Role of a Virtual Chassis" on page 129.

We recommend that you specify the same primary-role priority value for the primary and backup members. In this example, the highest possible primary-role priority has been assigned to two members. However, the member that was powered on first is given priority according to the primary-role election algorithm. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62. The other members use the default primary-role priority in this example, and they take on the linecard role.

NOTE: Use the `request virtual-chassis renumber` command if you want to change the member ID that the primary assigns to a member switch.

Configuring an EX4600 Virtual Chassis with a Preprovisioned Configuration File

When you preprovision a Virtual Chassis configuration, you assign the member ID and role for each switch in the Virtual Chassis.

To set up a Virtual Chassis using a preprovisioned configuration:

NOTE: You can configure a EX4600 Virtual Chassis while the cables are or are not physically connected.

1. Make a list of the serial numbers of all the switches to be connected in a Virtual Chassis configuration.
2. Note the desired role (`routing-engine` or `line-card`) of each switch. If you configure the member with a `routing-engine` role, it is eligible to function in the primary or backup role. If you configure the member with a `line-card` role, it is not eligible to function in the primary or backup role.

   If you are configuring a mixed EX4300 and EX4600 Virtual Chassis, EX4300 switches must be configured into the `line-card` role only. A mixed EX4300 and EX4600 Virtual Chassis
3. Power on only the switch that you plan to use as the primary switch.
4. (Required for a mixed Virtual Chassis only) Set the primary switch into mixed mode, and reboot the switch for the change to take effect:

   ```
   user@device> request virtual-chassis mode mixed reboot
   ```

5. If you are configuring a mixed Virtual Chassis, wait for the reboot to complete before performing this step.Specify the identification parameters for the switch by completing the initial configuration. See *Configuring Junos OS on the EX4600*.

   NOTE: The properties that you specify for the primary switch apply to the entire Virtual Chassis configuration.

6. (Optional) Configure the primary switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis:

```
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```

7. Specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@switch# set preprovisioned
```

8. Specify all the members that you want included in the Virtual Chassis, listing each switch's serial number with the desired member ID and role:

```
[edit virtual-chassis]
user@switch# set member 0 serial-number abc123 role routing-engine
user@switch# set member 1 serial-number def456 role routing-engine
user@switch# set member 2 serial-number ghi789 role line-card
user@switch# set member 3 serial-number jkl012 role line-card
```

9. (Optional. Recommended for a two-member Virtual Chassis) Disable the split and merge feature:

```
[edit virtual-chassis]
user@switch# set no-split-detection
```

   See "Disabling Split and Merge in a Virtual Chassis" on page 140 for more details.

10. Power on the other member switches. The member IDs and roles have been determined by the configuration, so you can power on the member switches in any order.

11. (Required if you are configuring a mixed Virtual Chassis) Set each individual switch into mixed mode, and reboot the switch for the change to take effect:

```
user@device> request virtual-chassis mode mixed reboot
```

12. On each individual member switch, configure the ports you will use to interconnect the member switches into VCPs using the following command:

```
user@switch> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
local
```

where *pic-slot-number* is the PIC slot number.

For instance, if you wanted to set port 0 on the QSFP+ interface on PIC slot 2 as a VCP:

```
user@switch> request virtual-chassis vc-port set pic-slot 2 port 0 local
```

The VCPs automatically bundle into a Link Aggregation Group when two or more interfaces of the same speed are configured into VCPs between the same two member switches. See "Understanding Virtual Chassis Port Link Aggregation" on page 68.

NOTE: You cannot modify the primary-role priority when you use a preprovisioned configuration. The primary-role priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two Routing Engines are assigned the same primary-role priority value. However, the member that was powered on first has higher prioritization according to the primary-role election algorithm. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62.

RELATED DOCUMENTATION

Configuring Primary Role of a Virtual Chassis | 129

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Configuring an EX4650 or a QFX Series Virtual Chassis

IN THIS SECTION

This topic discusses configuring an EX4650 or a QFX Series Virtual Chassis. For information on configuring a Virtual Chassis Fabric (VCF), see *Understanding Virtual Chassis Fabric Configuration*.

You configure a Virtual Chassis by configuring Virtual Chassis ports (VCPs) on the member switches and interconnecting the switches using the VCPs. The VCPs pass all data and control traffic between member switches in the Virtual Chassis. See "Understanding QFX Series Virtual Chassis" on page 25 and "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details on the switches that can be interconnected into a Virtual Chassis, and the ports on those switches that can be used as VCPs.

When you set up a Virtual Chassis, ideally all the proposed member switches should have the default factory configuration and operating in standalone mode.

- All the switches interconnected into a Virtual Chassis must be running the same version of Junos OS. See Installing Software Packages on QFX Series Devices.
- For QFX3500 or QFX3600 switches in a Virtual Chassis, you must download the Junos OS image that supports Virtual Chassis—an image that includes “jinstall-qfx-3-” in the filename when the Junos OS image is downloaded from the Software Center. QFX3500 and QFX3600 switches that are node devices in a QFabric system cannot be part of a Virtual Chassis.
- For QFX5100 and EX4300 switches in a QFX5100 Virtual Chassis, you must download the software image for the standalone switch. EX Series and QFX switches that are in a Junos Fusion cannot be part of a Virtual Chassis.
- For a QFX5110 Virtual Chassis with both QFX5110 and QFX5100 switches, all the switches must be running the same Junos OS image that includes “-qfx-5e-” in the Junos OS software package filename.

CAUTION: You *must* upgrade QFX5100 switches running a Junos OS image with "-qfx-5-" in the software package filename to a "-qfx-5e-" image filename before adding them to a QFX5110 Virtual Chassis, or the Virtual Chassis will not form. See *Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric*.

You can set up the Virtual Chassis with either:

- A preprovisioned configuration—With preprovisioning, you deterministically control the member ID and role assigned to a member switch by tying it to its serial number.
- A nonprovisioned configuration—Without provisioning, the primary sequentially assigns a member ID to other member switches, and determines the role of each member switch using the primary-role priority value and other factors in the primary-role election algorithm.

A Virtual Chassis configuration has two switches acting in the Routing Engine role—the primary switch and the backup switch. With any Virtual Chassis configuration, we recommend that you always use `commit synchronize` rather than simply `commit` to save configuration changes. This make sure the configuration changes are saved to both switches acting as Routing Engines.

Be sure that all switches that are interconnected into a Virtual Chassis are running the same version of Junos OS. See Installing Software Packages on QFX Series Devices.

Understanding the Licensing Requirements for a Virtual Chassis

Feature licenses are required to configure higher tier features on a Virtual Chassis. WIth Juniper Flex licensing, a license is required on all members of a Virtual Chassis.

For information on the feature licensing requirements for a Virtual Chassis, see Licenses for EX Series or Software Features That Require Licenses on the QFX Series.

You can install the feature licenses after configuring the Virtual Chassis.

Configuring an EX4650 or QFX Series Virtual Chassis with a Preprovisioned Configuration

Preprovisioning a Virtual Chassis configuration allows you to assign the member ID and role for each switch in the Virtual Chassis.

Before you begin, note that you can configure a Virtual Chassis while the cables are or are not physically connected. However, when committing a preprovisioned configuration on the member switches, under certain conditions the switches will *autoprovision* the ports cabling the member switches together, which means that those ports are automatically converted into Virtual Chassis ports (VCPs) when the Virtual Chassis members detect the link. Conditions for autoprovisioning include having LLDP enabled on the interfaces being used for the VCP links, and neither side of the link already has the port set as a VCP.

Automatic conversion of VCP links can cause links to come up unexpectedly, so if you want to control when the VCP links become active during Virtual Chassis configuration on the member switches, before you start the configuration, you can disable the VCP auto-conversion feature or any of the conditions required for the feature. See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45 for details. If the conditions for autoprovisioning the links are not present, you must manually set the ports connecting the member switches as VCPs as described in this procedure.

To set up a Virtual Chassis using a preprovisioned configuration:

1. Make a list of the serial numbers of all the switches to be connected in a Virtual Chassis configuration.

   > **NOTE**: Serial number values are case-sensitive.

2. Note the desired role (`routing-engine` or `line-card`) of each switch. If you configure the member with a `routing-engine` role, it is eligible to function as the primary or backup Routing Engine . If you configure the member with a `line-card` role, it is not eligible to function as the primary or backup. See "Understanding Virtual Chassis Components" on page 35 and "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details on which switches are recommended or required to be configured into each role depending on the types of switches in the Virtual Chassis.
3. Power on only the switch that you plan to use as the primary switch.
4. Specify the identification parameters for the switch by completing the initial configuration. See *Configuring Junos OS on the EX4650*, Configuring a QFX3500 Device as a Standalone Switch or Configuring a QFX3600 Device as a Standalone Switch, Configure a QFX5100 Device, Configuring a QFX5110, Configure Junos OS on the QFX5120, or Performing the Initial Software Configuration for QFX5200 Switches.

   > **NOTE**: The properties that you specify for the primary switch apply to the entire Virtual Chassis configuration.

5. (Optional) Configure the primary switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis:

```
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```

6. (Required for a mixed Virtual Chassis only) Set the primary switch into mixed mode and reboot the switch to complete the configuration:

   NOTE: You must complete this step if your Virtual Chassis includes a combination of different types of switches, except when you are configuring a Virtual Chassis that includes only QFX3500 and QFX3600 switches or only QFX5110 and QFX5100 switches (which are considered to be non-mixed Virtual Chassis). See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details.

```
user@device> request virtual-chassis mode mixed reboot
```

7. After the reboot is complete, specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@switch# set preprovisioned
```

8. Specify all the members that you want included in the Virtual Chassis, listing each switch's serial number with the desired member ID and role:

```
[edit virtual-chassis]
user@switch# set member 0 serial-number abc123 role routing-engine
user@switch# set member 1 serial-number def456 role routing-engine
user@switch# set member 2 serial-number ghi789 role line-card
user@switch# set member 3 serial-number jkl012 role line-card
```

9. (Optional. Recommended for a two-member Virtual Chassis) Disable the split and merge feature:

```
[edit virtual-chassis]
user@switch# set no-split-detection
```

10. Commit the configured items.

11. Power on the other member switches. The member IDs and roles have been determined by the configuration, so you can power on the member switches in any order.
12. (Required if you are configuring a mixed Virtual Chassis) Set each individual switch into mixed mode and reboot the switch to complete the configuration:

    NOTE: You must complete this step if your Virtual Chassis includes a combination of different types of switches, except when you are configuring a Virtual Chassis that includes only QFX3500 and QFX3600 switches or only QFX5110 and QFX5100 switches (which are considered to be non-mixed Virtual Chassis). See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details.

    ```
    user@device> request virtual-chassis mode mixed reboot
    ```

13. (Optional, if the VCP links will not be autoprovisioned) On each individual member switch, configure the ports that you will use to interconnect the member switches into VCPs:

    NOTE: SFP+, QSFP+, and 10-Gbps copper links between two member switches will be automatically set as VCPs (auto-provisioned) in a preprovisioned configuration under certain conditions. The required conditions on both sides of the link include having automatic VCP conversion enabled on the switches, LLDP enabled on the interfaces being used as VCPs, and those ports are not already set as VCPs. . (See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45.)

    This step is optional and should only be used when a VCP link is not automatically created.

    ```
    user@switch> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
    ```

    where *pic-slot-number* is the PIC slot number.

    For example, if you wanted to set port 0 on the QSFP+ interface on PIC slot 2 as a VCP:

    ```
    user@switch> request virtual-chassis vc-port set pic-slot 2 port 0
    ```

NOTE: You can include the `local` option if you want to make sure the command applies only to that port locally on the switch where you're running the command.

The VCPs automatically bundle into a Link Aggregation Group when two or more interfaces of the same speed are configured into VCPs between the same two member switches. See "Understanding Virtual Chassis Port Link Aggregation" on page 68.

NOTE: You cannot modify the primary-role priority when you are using a preprovisioned configuration. The primary-role priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two Routing Engines are assigned the same primary-role priority value. However, the member that was powered on first has higher priority according to the primary-role election algorithm. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62.

NOTE: If you want to change the member ID of a preprovisioned member switch later, you must re-configure the member information for that member switch using the `member` configuration statement. For example, to re-configure member 3 (which has serial number “jkl012”) to have 6 as its new member ID, you associate member ID 6 with the serial number of member 3, and then delete the configuration item for member ID 3, as follows:

```
[edit virtual-chassis]
user@switch# set member 6 serial-number jkl012
user@switch# delete member 3
```

The `request virtual-chassis renumber` command can only be used to change a Virtual Chassis member ID in a nonprovisioned Virtual Chassis.

Configuring an EX4650 or a QFX Series Virtual Chassis with a Nonprovisioned Configuration

You can use a nonprovisioned configuration to set up an EX4650 or a QFX Series Virtual Chassis.

To configure the Virtual Chassis using a nonprovisioned configuration:

NOTE: You can configure the Virtual Chassis while the cables are or are not physically connected.

1. Power on only the switch that you plan to use as the primary switch.
2. (Required for a mixed Virtual Chassis only) Set the primary switch into mixed mode and reboot the switch to complete the configuration:

   NOTE: You must complete this step if your Virtual Chassis includes a combination of different types of switches, except when you are configuring a Virtual Chassis that includes only QFX3500 and QFX3600 switches or only QFX5110 and QFX5100 switches (which are considered to be non-mixed Virtual Chassis). See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details.

   ```
   user@device> request virtual-chassis mode mixed reboot
   ```

3. After the primary switch reboots, specify the identification parameters for the switch by completing the initial configuration. See *Configuring Junos OS on the EX4650*, Configuring a QFX3500 Device as a Standalone Switch, Configuring a QFX3600 Device as a Standalone Switch, Configure a QFX5100 Device, Configuring a QFX5110, Configure Junos OS on the QFX5120, or Performing the Initial Software Configuration for QFX5200 Switches for details.

   NOTE: The properties that you specify for the primary switch apply to the entire Virtual Chassis configuration.

4. (Optional) Configure the primary switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis:

   ```
   user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
   ```

5. (Optional) Configure primary-role priority for the member switches:

   ```
   [edit virtual-chassis]
   user@switch# set member 0 mastership-priority 255
   user@switch# set member 1 mastership-priority 255
   ```

6. (Optional. Recommended for a two-member Virtual Chassis) On the primary switch, disable the split and merge feature:

```
[edit virtual-chassis]
user@switch# set no-split-detection
```

7. Commit the configured items.
8. Power on the other member switches.
9. (Required for a mixed Virtual Chassis only) Set each individual switch into mixed mode and reboot the switch to complete the configuration:

   NOTE: You must complete this step if your Virtual Chassis includes a combination of different types of switches, except when you are configuring a Virtual Chassis that includes only QFX3500 and QFX3600 switches or only QFX5110 and QFX5100 switches (which are considered to be non-mixed Virtual Chassis). See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details.

```
user@device> request virtual-chassis mode mixed reboot
```

10. On each individual member switch, configure the ports that will be used to interconnect the member switches into VCPs using the following command:

```
user@switch> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
```

   where *pic-slot-number* is the PIC slot number.

   For example, if you wanted to set port 0 on the QSFP+ interface on PIC slot 2 as a VCP:

```
user@switch> request virtual-chassis vc-port set pic-slot 2 port 0
```

   NOTE: You can include the `local` option if you want to make sure the command applies only to that port locally on the switch where you're running the command.

The VCPs automatically bundle into a Link Aggregation Group when two or more interfaces of the same speed are configured into VCPs between the same two member switches. See "Understanding Virtual Chassis Port Link Aggregation" on page 68.

NOTE: If you don't set primary role priorities on any member switches, the default primary-role priority value is`128`. If the primary role priorities are the same on all members, the primary-role election algorithm determines the member switches in the primary and backup Routing Engine roles based on the other factors in the algorithm. You control the role settings by configuring the primary-role priority to a higher number on the switches you want to be the primary and backup members. .(see "Configuring Primary Role of a Virtual Chassis" on page 129). We recommend that you specify the same primary-role priority value for the members you want to be the primary and backup members. This example assigns the highest possible primary-role priority to two members. However, the member that was powered on first has higher priority according to the primary-role election algorithm. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62 for details on all the factors considered when electing the primary. The other members have the default primary-role priority in this example, and they become linecard role members.

NOTE: If you want to change the member ID that the primary automatically assigned to a member switch, use the `request virtual-chassis renumber` command.

RELATED DOCUMENTATION

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Adding a New Switch to an Existing EX2300, EX3400, EX4300, or EX4400 Virtual Chassis

You can use this procedure to:

- Add an EX2300 switch to an existing EX2300 Virtual Chassis.

NOTE: Starting with Junos OS Release 18.4R1, you can combine any models of EX2300 switches, including EX2300 multigigabit models, into the same EX2300 Virtual Chassis using this procedure. In releases prior to Junos OS Release 18.4R1, EX2300 multigigabit switches cannot be combined with any other models of EX2300 switches in the same Virtual Chassis.

- Add an EX3400 switch to an existing EX3400 Virtual Chassis.
- Add an EX4300 switch to an existing non-mixed EX4300 Virtual Chassis. For example:
  - Add an EX4300 multigigabit model (EX4300-48MP) switch to a Virtual Chassis consisting of all EX4300 multigigabit model switches.
  - Add any other type of EX4300 switch to a Virtual Chassis that contains only EX4300 switches and no EX4300 multigigabit model switches.
- Add an EX4300 switch to a mixed EX4300 Virtual Chassis that consists of a supported combination of EX4300 multigigabit model switches and any other EX4300 switches.
- Add any model EX4400 switch (including EX4400 multigigabit models) to an existing EX4400 Virtual Chassis.

You can't use this procedure to:

- Add an EX4300 multigigabit model (EX4300-48MP) switch to an existing EX4300 Virtual Chassis that consists only of other non-multigigabit EX4300 model switches.

  EX4300 multigigabit model switches must be in the primary and backup Routing Engine roles in an EX4300 Virtual Chassis that contains both types of switches. As a result, in that case, you should first create a new multigigabit EX4300 Virtual Chassis with multigigabit EX4300 switches in the primary and backup roles, and then merge the other non-multigigabit EX4300 model switches from the original non-mixed Virtual Chassis into a new mixed-mode EX4300 Virtual Chassis.

- Add an EX4300 switch (non-multigigabit models) to a supported mixed Virtual Chassis or VCF.

  EX4300 switches that are not multigigibit model can be part of a mixed Virtual Chassis with EX4600 switches or a mixed Virtual Chassis or VCF with particular QFX Series switches. See these other references for how to add an EX4300 switch in those cases:

  - Adding an EX4300 switch to a mixed Virtual Chassis with EX4600 switches: "Adding an EX4600 Switch to a Mixed or Non-mixed Virtual Chassis" on page 113.
  - Adding an EX4300 switch to a mixed Virtual Chassis with QFX Series switches: "Adding a New Switch to an Existing EX4650 or QFX Series Virtual Chassis" on page 116
  - Adding an EX4300 switch to a mixed VCF: *Adding a Device to a Virtual Chassis Fabric*.

Before you begin, be sure you have:

- Confirmed that the new switch is supported as a member of the Virtual Chassis and in the role in which you want to add it. See "Understanding EX Series Virtual Chassis" on page 9 and "Understanding Virtual Chassis Components" on page 35 for details on the different EX Series switches, switch combinations, and switch roles that are supported or recommended in a Virtual Chassis.

- Ensured the new switch has the same version of Junos OS that is running on the Virtual Chassis primary switch, or the existing Virtual Chassis has the automatic software update feature configured.

  If you have configured the automatic software update feature in the existing Virtual Chassis, the primary switch updates newly added members with the correct software version automatically, if necessary. Otherwise, you must manually install the correct software version. See "Understanding Automatic Software Update on Virtual Chassis Member Switches" on page 74.

- Mounted the new switch in a rack.

- Determined which ports you will use as Virtual Chassis ports on the new switch, and the member ports in the existing Virtual Chassis to which you will interconnect the new switch.

  See "Virtual Chassis Port Options" on page 41 for the ports that can be used as VCPs on switches that support Virtual Chassis.

- If you are expanding a preprovisioned configuration:

  - Made a note of the serial number (the number is on the back of the switch). You will need to edit the Virtual Chassis configuration to include the serial number of the new member switch.

    NOTE: Serial number values are case-sensitive.

  - Edited the existing Virtual Chassis configuration to include the serial number of the new member switch. The parameters specified in the primary Virtual Chassis configuration file are applied to the new switch after it has been interconnected to an existing member switch.

  - (If you are using the autoprovisioning feature to add a member switch to an existing preprovisioned Virtual Chassis) Confirmed that the member ports in the Virtual Chassis to which you will interconnect the new switch are not already configured as VCPs. One condition for automatic VCP conversion is that the ports on both sides of the new link must not already be configured as VCPs. See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45 for details on the requirements and conditions under which this feature will be invoked.

- (Optional) Configured Ethernet interfaces on different member switches into the same LAG. For example, see *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*.

An active member switch might temporarily go down before coming back up as part of this procedure. Having traffic load-balanced across member switches using a LAG helps alleviate traffic loss during this procedure.

- Deleted the `no-split-detection` configuration item if you are expanding a two-member Virtual Chassis with this option configured. We very strongly recommend that you enable the split detection and merge feature for Virtual Chassis configurations with more than two members. This feature is enabled in the default configuration when you initially set up a Virtual Chassis.

To add a new member switch to an existing Virtual Chassis configuration:

1. If the new member switch has been previously configured, revert that switch's configuration to the factory defaults before interconnecting it into the Virtual Chassis. See *Reverting to the Default Factory Configuration for the EX Series Switch*.
2. (Required for a mixed EX4300 Virtual Chassis only) A mixed EX4300 Virtual Chassis contains a combination of EX4300 multigigabit model (EX4300-48MP) switches and other EX4300 switches.

   If you are adding a new switch in this case, set the new switch into mixed mode, and reboot the switch for the mode change to take effect as follows:

   - If the new switch is an EX4300 multigigabit model (EX4300-48MP) switch:

     ```
     user@device> request virtual-chassis mode mixed reboot
     ```

   - If the new switch is any other EX4300 model switch, you must also configure the switch with a special port mode by including the `ieee-clause-82` option when you set mixed mode. This port mode enables VCPs on the EX4300 switch to communicate with VCPs on multigigabit model member switches:

     ```
     user@device> request virtual-chassis mode ieee-clause-82 mixed reboot
     ```

   If adding this switch to the Virtual Chassis converts a non-mixed EX4300 Virtual Chassis into a mixed EX4300 Virtual Chassis, log into the Virtual Chassis and set all the existing member switches into mixed mode as well, and reboot the Virtual Chassis to complete the configuration:

   ```
   user@vc-master> request virtual-chassis mode mixed reboot all-members
   ```

   For example, if you have an EX4300 Virtual Chassis consisting of all EX4300 multigigabit model member switches, that is a non-mixed EX4300 Virtual Chassis. If you add an EX4300 switch that isn't a multigigabit model to that Virtual Chassis, you must change the mode to `mixed` on all existing members when you add the new switch.

Wait for the reboot operation to complete on any affected switches in this step before moving to the next step.

3. Power off the new switch, and interconnect the unpowered new switch to one member of the existing Virtual Chassis configuration using a supported VCP.

   Connect only one VCP on the unpowered new switch to a VCP on a member switch in the existing Virtual Chassis at this point in the procedure.

4. Power on the new switch.
5. Set the interconnecting ports as Virtual Chassis Ports (VCPs), if needed:

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port port-number
```

   You do not need to perform this step in the following cases:

   - You are using dedicated VCPs, which do not need to be configured. EX4300 multigigabit model (EX4300-48MP) switches have dedicated VCPs on the rear panel, which are the only VCP port options on those switches.

   - You do not typically need to perform this step with QSFP+ ports on EX3400 and EX4300 switches that are not multigigabit models. All QSFP+ ports on these switches are configured as VCPs by default. However, you might use the `request virtual-chassis vc-port` command to set a QSFP+ port back into a VCP if the QSFP+ port had previously been reconfigured as a network port.

   - The default VCPs on EX4400 switches are the only ports that can be used as VCPs, so you must use those ports to interconnect member switches in an EX4400 Virtual Chassis. However, if you previously converted the default VCPs into network ports by enabling network port mode on the switch, you must disable network port mode and reboot the switch to convert them back into VCPs. Use the `request virtual-chassis mode network-port disable <reboot>` command to do this if needed. You can enter the `show virtual-chassis mode` command to see whether network port mode is enabled.on the switch.

   - If you have the right conditions to use the autoprovisioning feature, after the new switch is provisioned and cabled into the Virtual Chassis, the interconnecting links automatically convert into VCP links. You do not need to manually set the ports on either side of the links as VCPs.

6. Confirm that the new member switch is now included within the Virtual Chassis configuration by entering the `show virtual-chassis` command. The new member switch should be listed in the output and the `Status` is `Prsnt`.
7. Cable the next port into the Virtual Chassis, configuring the ports into VCPs if needed.

CAUTION: If you immediately cable both VCPs on the new switch into the existing Virtual Chassis at the same time, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. Network traffic to this switch is dropped during the downtime.

The member switch will return to the normal operational state with no user intervention, and normal operation of the Virtual Chassis will resume after this downtime.

RELATED DOCUMENTATION

Configuring an EX2300, EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis | 80

Understanding Virtual Chassis Components | 35

Adding an EX4600 Switch to a Mixed or Non-mixed Virtual Chassis

You can use this procedure to add an EX4600 switch to a mixed or non-mixed Virtual Chassis or to add an EX4300 switch to an existing mixed EX4300 and EX4600 Virtual Chassis. For EX4650 Virtual Chassis, which is more like a QFX5120 Virtual Chassis than an QFX4600 Virtual Chassis, see "Adding a New Switch to an Existing EX4650 or QFX Series Virtual Chassis" on page 116.

NOTE: EX4300 multigigabit model (EX4300-48MP) switches are not supported in a mixed Virtual Chassis with EX4600 switches.

Before you begin, be sure you have:

- Mounted the new switch in a rack.
- Determined which ports you will use as Virtual Chassis ports on the new switch, and the member ports in the existing Virtual Chassis to which you will interconnect the new switch.
- If you are expanding a preprovisioned configuration:
  - Made a note of the serial number (the number is on the back of the switch). You will need to edit the Virtual Chassis configuration to include the serial number of the new member switch.

> **NOTE**: Serial number values are case-sensitive.

- Edited the existing Virtual Chassis configuration to include the serial number of the new member switch. The parameters specified in the primary Virtual Chassis configuration file are applied to the new switch after it has been interconnected to an existing member switch.

- (If you are using the autoprovisioning feature to add a member switch to an existing preprovisioned Virtual Chassis) Confirmed that the member ports in the Virtual Chassis to which you will interconnect the new switch are not already configured as VCPs. One condition for automatic VCP conversion is that the ports on both sides of the new link must not already be configured as VCPs. See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45 for details.

- (Optional) Configured Ethernet interfaces on different member switches into the same LAG. See *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*.

  An active member switch might temporarily go down before coming back up as part of this procedure. Having traffic load-balanced across member switches using a LAG helps alleviate traffic loss during this procedure.

To add a new member switch to an existing Virtual Chassis configuration:

1. If the new member switch has been previously configured, revert that switch's configuration to the factory defaults before interconnecting it into the Virtual Chassis. See *Reverting to the Default Factory Configuration for the EX Series Switch*.
2. (Required for a mixed Virtual Chassis) Set the new switch into mixed mode, and reboot the switch to complete the configuration:

   ```
   user@device> request virtual-chassis mode mixed reboot
   ```

   If adding this switch to the Virtual Chassis converts a non-mixed Virtual Chassis into a mixed Virtual Chassis, log into the Virtual Chassis and set the switches into mixed mode. Reboot the Virtual Chassis to complete the configuration:

   ```
   user@device> request virtual-chassis mode mixed reboot all-members
   ```

3. If you are rebooting the switch or the Virtual Chassis to complete a mixed mode setting change, wait for the reboot to complete before performing this step.Power off the new switch, and interconnect the unpowered new switch to one member of the existing Virtual Chassis configuration using a port that is supported as a VCP.

Connect only one VCP on the unpowered new switch to a VCP on a member switch in the existing Virtual Chassis at this point of the procedure.

4. Set the interconnecting QSFP+ or SFP+ ports as Virtual Chassis ports (VCPs) on the new member switch and the existing Virtual Chassis member switch where you connected the new switch, if needed:

```
user@switch> request virtual-chassis vc-port set pic-slot pic-slot port port-number
```

NOTE: Include the `local` option in this command if you want to make sure the command applies only to that port locally on the switch where you're running the command.

You do not need to perform this step in a preprovisioned Virtual Chassis if you set up the right conditions to use the autoprovisioning feature (see "Automatic Virtual Chassis Port (VCP) Conversion" on page 45). After the new switch is provisioned and cabled into the Virtual Chassis, the interconnecting links automatically convert into VCP links. You do not need to manually set the ports on either side of the links as VCPs.

5. Confirm that the new member switch is now included within the Virtual Chassis configuration by entering the `show virtual-chassis` command. The new member switch should be listed in the output and the `Status` is `Prsnt`.

6. Cable the next port into the Virtual Chassis, using Steps 2 through 5.

CAUTION: If you immediately cable both VCPs on the new switch into the existing Virtual Chassis at the same time, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. Network traffic to this switch is dropped during the downtime.

The member switch will return to the normal operational state with no user intervention, and normal operation of the Virtual Chassis will resume after this downtime.

RELATED DOCUMENTATION

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis | 93

Understanding Virtual Chassis Components | 35

Adding a New Switch to an Existing EX4650 or QFX Series Virtual Chassis

IN THIS SECTION

Use this procedure to add a new switch to an EX4650 or a QFX Series Virtual Chassis.

To add a switch to a Virtual Chassis Fabric (VCF), see *Adding a Device to a Virtual Chassis Fabric*.

Add a New Switch to an EX4650 or QFX Series Virtual Chassis

You can use this procedure to add a switch in a supported combination to an existing EX4650 or QFX Series Virtual Chassis. A Virtual Chassis is a supported combination of switches interconnected using Virtual Chassis ports (VCPs). EX4650, QFX5120 and QFX5200 switches can be members of a non-mixed Virtual Chassis only (all members must be the same type of switch). QFX series switches that can be members of a mixed or non-mixed Virtual Chassis include QFX3500, QFX3600, QFX5100, and QFX5110 switches.,. EX4300 switches can also be members of a mixed QFX Series Virtual Chassis with QFX3500, QFX3600, and QFX5100 switches. See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54 for details on the types and combinations of switches that can make up a mixed QFX Series Virtual Chassis.

NOTE: EX4300 multigigabit model (EX4300-48MP) switches are not supported in a mixed Virtual Chassis with QFX Series switches.

Before you begin, make sure you have:

- Mounted the new switch in a rack.
- Enabled automatic software updates on the Virtual Chassis. See "Configuring Automatic Software Update on Virtual Chassis Member Switches" on page 141.

CAUTION: When adding a QFX5100 switch installed with a “-qfx-5-” Junos OS image to a QFX5110 Virtual Chassis, you must first upgrade the QFX5100 switch to run a “-qfx-5e-” Junos OS image.before this step in the procedure. The automatic software update feature cannot automatically update a QFX5100 switch running a “-qfx-5-” Junos OS image to a “-qfx-5e-” image. See *Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric*

- Determined which ports you will use as Virtual Chassis ports on the new switch, and the member ports in the existing Virtual Chassis to which you will interconnect the new switch.
- If you are expanding a preprovisioned configuration:
  - Made a note of the serial number (the number is on the back of the switch). You’ll need to edit the Virtual Chassis configuration to include the serial number of the new member switch.

    NOTE: Serial number values are case-sensitive.

  - Edited the existing Virtual Chassis configuration to include the serial number of the new member switch. The Virtual Chassis applies the parameters in the primary’s configuration file to the new switch after it has been interconnected with an existing member switch.
  - (If you are using the autoprovisioning feature to add a member switch to an existing preprovisioned Virtual Chassis) Confirmed that the member ports in the Virtual Chassis to which you will interconnect the new switch are not already configured as VCPs. One condition for automatic VCP conversion is that the ports on both sides of the new link must not already be configured as VCPs. See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45 for details.
- (Optional) Configured Ethernet interfaces on different member switches into the same LAG. See Configuring Link Aggregation.

  An active member switch might temporarily go down before coming back up as part of this procedure. If you configure the Virtual Chassis members so that traffic is load-balanced across member switches using a LAG, you can alleviate traffic loss during this procedure.

To add a new member switch to an existing EX4650 or QFX Series Virtual Chassis configuration:

1. If you previously configured the new member switch, we recommend you revert that switch’s configuration to the factory defaults before interconnecting it into the Virtual Chassis. See *Reverting to the Default Factory Configuration*.
2. (Recommended for a QFX5100 Virtual Chassis under certain conditions) When you add or replace a QFX5100-24Q switch that is configured in the Routing Engine role in a QFX5100 Virtual Chassis, if the new switch has two EX4600-EM-8F expansion modules, we recommend that you set the primary

role priorities on the routing engine members and the new switch to prevent a primary-role switchover to the new switch until after the new switch is completely initialized in the Virtual Chassis.

Before interconnecting the new switch into the Virtual Chassis in this case, see "Add or Replace a QFX5100-24Q Switch with Two Expansion Modules in a QFX5100 Virtual Chassis" on page 120 for details on why, when, and how you should do this step.

> **NOTE**: You might need to do this even if the new switch has the default factory configuration.

3. (Required for a mixed Virtual Chassis only) Set the new switch into mixed mode and reboot the switch to complete the configuration:

> **NOTE**: You do not need to configure your Virtual Chassis into mixed mode if the Virtual Chassis is composed of only QFX3500 and QFX3600 switches or only QFX5110 and QFX5100 switches (which are considered non-mixed Virtual Chassis).

```
user@device> request virtual-chassis mode mixed reboot
```

If you are adding a switch that converts a non-mixed Virtual Chassis into a mixed Virtual Chassis, you must also log onto the Virtual Chassis and enter the **request virtual-chassis mode mixed all-members reboot** command either before or after interconnecting the new switch into your Virtual Chassis. This step makes sure all switches in the Virtual Chassis can communicate with the new mixed-mode member switch.

4. Interconnect the new switch to one member of the existing Virtual Chassis using an interface that can be configured into a VCP. See "Virtual Chassis Port Options" on page 41 for details on ports you can use as VCPs on different switches.

   Connect only one interface on the new switch to a VCP on a member switch in the existing Virtual Chassis at this point of the procedure.

5. Set the interconnecting ports for the new member switch as Virtual Chassis Ports (VCPs) on the new member switch and the existing Virtual Chassis member switch where you connected the new switch, if needed:

```
user@switch> request virtual-chassis vc-port set pic-slot slot-number port port-number
```

NOTE: Include the `local` option in this command if you want to make sure the command applies only to that port locally on the switch where you're running the command.

You do not need to perform this step in a preprovisioned Virtual Chassis if you set up the right conditions to use the autoprovisioning feature (see "Automatic Virtual Chassis Port (VCP) Conversion" on page 45). After the new switch is provisioned and cabled into the Virtual Chassis, the interconnecting links automatically convert into VCP links. You do not need to manually set the ports on either side of the links as VCPs.

6. Confirm that the new member switch is now included within the Virtual Chassis configuration by entering the `show virtual-chassis` command on a Virtual Chassis console or management port. You should see the new member switch listed in the output with `Status` displayed as `Prsnt`.
7. Cable the next port into the Virtual Chassis. Refer to Steps 4 through 6 in this procedure.

CAUTION: If you immediately cable both VCPs on the new switch into the existing Virtual Chassis at the same time, a member switch that was already part of the Virtual Chassis might become non-operational for several seconds. The Virtual Chassis drops network traffic to this switch during the downtime.

The member switch returns to the normal operational state with no user intervention, and normal operation of the Virtual Chassis resumes after this downtime.

8. Split detection is enabled in a Virtual Chassis by default. We recommend configuring the `no-split-detection` option to disable it in a Virtual Chassis that has only two members. We very strongly recommend keeping it enabled in a Virtual Chassis that has more than two members. See "Understanding Split and Merge in a Virtual Chassis" on page 71 for details.

   If your Virtual Chassis had only two members with `no-split-detection` configured, now that you've added another member switch in this procedure, enable split detection again (in other words, remove the `no-split-detection` setting):

```
[edit]
user@switch# delete virtual-chassis no-split-detection
```

9. If you need to customize your Virtual Chassis configuration further, see "Configuring an EX4650 or a QFX Series Virtual Chassis" on page 100 or *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*.

SEE ALSO

Configuring an EX4650 or a QFX Series Virtual Chassis | 100
Understanding Virtual Chassis Components | 35

Add or Replace a QFX5100-24Q Switch with Two Expansion Modules in a QFX5100 Virtual Chassis

A QFX5100-24Q switch installed with two EX4600-EM-8F expansion modules needs more time to check and initialize all its interfaces than a switch with one or no expansion modules. If you add a new switch or replace a member switch in an QFX5100 Virtual Chassis, when you connect and power on the new switch, the Virtual Chassis detects the new switch and reconfigures the Virtual Chassis to include it.

During that time, the Virtual Chassis might reassign primary role to the new switch before it is fully initialized and ready to take the primary Routing Engine role. This might happen if the new switch

- is a QFX5100 switch with two expansion modules,
- is configured in the Routing Engine role,
- has a system MAC address lower than the other member switches in the Routing Engine role, and
- has the same primary-role priority as the other member switches in the Routing Engine role.

See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62 for more information on primary-role election.

The Virtual Chassis takes significantly longer to re-stabilize across a primary-role switchover during Virtual Chassis reconfiguration, which causes traffic loss. To avoid a primary-role switchover in this case, before cabling the new switch into the Virtual Chassis, you can reassign the new switch’s primary-role priority to be lower than the other Virtual Chassis member switches configured in the Routing Engine role.

For example, you have a QFX5100 Virtual Chassis with member 0 as the primary routing engine and member 1 as the backup routing engine. Both members have primary-role priority 255 in the current Virtual Chassis configuration, and you want to replace member 0 with a QFX5100-24Q switch that has two expansion modules.

1. Remove member 0 according to the procedure in "Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch" on page 123.

   Member 1 assumes primary role of the Virtual Chassis.

2. Reconfigure the primary-role priority of member 0 on the Virtual Chassis to a lower number than member 1 to make sure member 1 retains primary role while you replace member 0:

```
[edit]
user@switch# set virtual-chassis member 0 mastership-priority 128
user@switch# set virtual-chassis member 1 mastership-priority 255
user@switch# commit synchronize
```

3. Add the new member 0 into the Virtual Chassis using the procedure in "Add a New Switch to an EX4650 or QFX Series Virtual Chassis" on page 116, or continue with the steps in "Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration" on page 125 after removing the member you are replacing.
4. After the new member switch is fully initialized and established as a member of the Virtual Chassis, you can restore the primary role priorities on member 0 so the Virtual Chassis switches primary role back to the new member 0:

```
[edit]
user@switch# set virtual-chassis member 0 mastership-priority 255
user@switch# commit synchronize
```

See "Configuring Primary Role of a Virtual Chassis" on page 129 for more details on changing primary role priorities.

RELATED DOCUMENTATION

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

IN THIS SECTION

You can remove or replace a member switch in a Virtual Chassis without disrupting network service on the other member switches.

If you remove a member switch, you can free up the member ID so it is available to be assigned to a new member switch later.

When you add a new member switch, the Virtual Chassis assigns the next available member ID to it. The Virtual Chassis retains the existing configuration items specific to particular member IDs. The Virtual Chassis applies those items to a replacement member switch that has the same member ID. By default, the Virtual Chassis applies configuration items that are not member-specific to all member switches.

NOTE: When you add or delete member switches in a Virtual Chassis configuration, internal routing changes might cause a sub-second temporary traffic loss. Also, if removing a member switch changes a mixed Virtual Chassis into a non-mixed Virtual Chassis, you must remove the mixed mode setting on all member switches of the Virtual Chassis and reboot the Virtual Chassis; network services are disrupted until the Virtual Chassis is up again.

This topic does not apply to:

- A Virtual Chassis Fabric (VCF).

  Instead, see Removing a Device From a Virtual Chassis Fabric for VCF information.

- A mixed Virtual Chassis that contains EX4200, EX4500, or EX4550 switches.

  Instead, see Removing an EX4200, EX4500, or EX4550 Switch From a Mixed Virtual Chassis (CLI Procedure).

- An EX8200 Virtual Chassis.

To remove or replace a member switch of any other EX Series or QFX Series Virtual Chassis, use one of the following procedures that matches what you want to do.

Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch

To remove a switch from a Virtual Chassis without replacing it:

1. Power off and disconnect the member switch you want to remove from the Virtual Chassis.
2. If the Virtual Chassis configuration is preprovisioned, on the Virtual Chassis primary, remove the removed switch's member setting from the preprovisioned configuration.

   ```
   [edit virtual-chassis]
   user@vc-primary# delete member removed-member-id
   ```

   If the Virtual Chassis configuration is nonprovisioned, change the `mastership-priority` values of each member switch as needed to reconfigure the Virtual Chassis roles. See Configuring Primary Role of a Virtual Chassis.
3. (Optional) If removing a member switch leaves only two remaining member switches in the Virtual Chassis, we recommend you disable split detection in a two-member Virtual Chassis. See Disabling Split and Merge in a Virtual Chassis for details.
4. Commit any configuration changes made in previous steps.
5. (For a mixed Virtual Chassis only) If removing this member switch changes the Virtual Chassis from a mixed to a non-mixed Virtual Chassis, you must also remove the mixed-mode setting from the Virtual Chassis. See Understanding Mixed EX Series and QFX Series Virtual Chassis for details on the combinations of switches that comprise a mixed Virtual Chassis. You must reboot the Virtual Chassis for the mode change to take effect. To do this, you can include the `reboot` option with the `request virtual-chassis mode mixed disable` command that turns off mixed mode, or reboot all member switches of the Virtual Chassis separately when ready to do so, as shown below.

   ```
   user@vc-primary> request virtual-chassis mode mixed disable all-members
   user@vc-primary> request system reboot all-members
   ```

   NOTE: Step 7 describes how to remove the mixed mode and other settings from the removed switch if needed.

6. When you remove a member switch from a Virtual Chassis configuration, the primary keeps that member switch's member ID in reserve. Use the following command on the Virtual Chassis primary to make that member ID available for reassignment:

```
user@vc-primary> request virtual-chassis recycle member-id member-id
```

7. If you want to use the removed switch as a standalone switch, you must remove any Virtual Chassis configuration items and settings on that switch. For a smooth transition to a new role as a standalone switch, we recommend to revert the switch to its default factory configuration using the `request system zeroize` command, and then apply the configuration items you want on the switch.

   If you do not want to revert to default factory settings, use commands such as the following to remove Virtual Chassis settings for the mode and VCPs from the removed switch:

   a. If you removed the switch from a mixed-mode Virtual Chassis that is not an EX4300 mixed Virtual Chassis, disable the mixed-mode setting on the switch as follows:

   ```
   user@switch> request virtual-chassis mode mixed disable
   ```

   If you removed a non-multigigabit model EX4300 switch from a mixed EX4300 Virtual Chassis that contains multigigabit EX4300 switches (EX4300-48MP), when you disable mixed mode, you must also disable the special `ieee-clause-82` port mode on the removed switch if you want to reconfigure it as a standalone switch or use it in any other type of mixed Virtual Chassis or non-mixed Virtual Chassis. Otherwise, the VCPs on the switch will not connect with other Virtual Chassis members or those ports will not operate properly as network ports.

   In this case, to disable mixed mode and the port mode on the switch:

   ```
   user@switch> request virtual-chassis mode mixed ieee-clause-82 disable
   ```

   See Understanding EX4300 Multigigabit and Other EX4300 Model Switches in a Mixed EX4300 Virtual Chassis for more information about this special port mode on EX4300 switches.

   b. Delete the VCP settings for any ports that were used as VCPs:

   ```
   user@switch> request virtual-chassis vc-port delete pic-slot pic-slot port port-number
   ```

   c. Reboot the standalone switch for settings such as mode changes to take effect.

Remove, Repair, and Reinstall the Same Switch

If you need to repair a member switch, you can remove it from the Virtual Chassis configuration without disrupting network service for the other member switches. The primary stores the configuration for the member ID so that it can be reapplied when the member switch (with the same base MAC address) is reconnected.

To remove, repair, and reinstall the member switch:

1. Power off and disconnect the member switch to be repaired.
2. Repair, as necessary.
3. Reconnect the switch and power it on.

Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration

If you are unable to repair a member switch, you can replace it with a different member switch of the same type while retaining the previous configuration. The primary stores the configuration of the member switch that was removed. When you connect a different member switch, the primary assigns a new member ID, but the old configuration is still stored under the previous member ID of the previous member switch.

To remove and replace a switch and reapply the old configuration:

1. Power off and disconnect the member switch to be replaced.

   > **NOTE**: See Step 7 in "Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch" on page 123 for information on how to disable Virtual Chassis settings from the removed switch if you want to use that switch in a different configuration.

2. If the replacement member switch has been previously configured, revert that switch's configuration to the factory defaults. See the `request system zeroize` command.

   The replacement member switch should be powered on and running with the factory default configuration at the end of this step.

3. (Recommended for a QFX5100 Virtual Chassis under certain conditions) When you add or replace a QFX5100-24Q switch that is configured in the Routing Engine role in a QFX5100 Virtual Chassis, if the new switch has two EX4600-EM-8F expansion modules, we recommend that you set the primary role priorities on the routing engine members and the new switch to prevent a primary-role

switchover to the new switch until after the new switch is completely initialized in the Virtual Chassis.

Before interconnecting the new switch into the Virtual Chassis in this case, see Add or Replace a QFX5100-24Q Switch with Two Expansion Modules in a QFX5100 Virtual Chassis for details on why, when, and how you should do this step.

NOTE: You might need to do this even if the new switch has the default factory configuration.

4. (Required when automatic software update is not enabled on the Virtual Chassis and the new member switch is running a version of Junos OS that is different than the version of Junos OS running on the Virtual Chassis) Ensure that the correct version of Junos OS is or will be installed on the new member switch by performing *either* of the following tasks:

   - Enable automatic software update on the Virtual Chassis. See Configuring Automatic Software Update on Virtual Chassis Member Switches. The Virtual Chassis will automatically update the software on the replacement switch in a later step when it is interconnected and recognized as part of the Virtual Chassis. The replacement switch does not require any action in this case for this step.

   - Install the version of Junos OS running on the Virtual Chassis onto the new member switch before interconnecting it into the Virtual Chassis. For EX series switches, see Installing Software on an EX Series Switch with a Virtual Chassis or Single Routing Engine (CLI Procedure), or for QFX Series switches, see Software Installation and Upgrade Overview and Installing Software Packages on QFX Series Devices. In this case, at the end of this step, the replacement switch will be running with the new version of the software and should have the factory default configuration.

     CAUTION: You can only set up a QFX5110 Virtual Chassis with both QFX5110 and QFX5100 switches if they are running the same Junos OS image that includes "-qfx-5e-" in the software package filename (from the Junos OS Software Center). If the switch you are replacing in a QFX5110 Virtual Chassis is a QFX5100 switch that you previously installed with a "-qfx-5-" Junos OS image file, you *must* upgrade the replacement switch to a "-qfx-5e-" image instead before inserting it into the QFX5110 Virtual Chassis. See Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric. The automatic software update feature can't update a "-qfx-5-" image to a "-qfx-5e-" image.

5. Connect one link from the replacement member switch to the Virtual Chassis as follows, depending on which type of ports you are using:

   - If you are interconnecting a switch using dedicated Virtual Chassis Ports (VCPs), connect one dedicated VCP on the replacement member switch to a dedicated VCP on another member switch in the Virtual Chassis.

- If you are interconnecting a switch using optical ports configured as VCPs:

  On the replacement switch, configure the optical ports that you are using to connect to the Virtual Chassis as VCPs. (You should also configure the optical ports on the existing member switches in the Virtual Chassis where the replacement member switch will be connected, if they are not already configured.) To configure an optical port as a VCP:

```
user@switch> request
virtual-chassis vc-port set pic-slot 1 port port-number
```

  Connect one configured optical port VCP on the replacement switch to a configured optical port VCP on another member switch in the Virtual Chassis.

  NOTE: You can set optical port VCPs on a standalone switch before interconnecting one link into an existing Virtual Chassis, or set them after connecting the link. In either case, you must set the ports as VCPs for the primary to detect and complete the process of adding the switch as a member. For more information on setting up VCPs on EX Series switches, see Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port, and for details on which ports can be configured as VCPs in a QFX Series Virtual Chassis, see Understanding Virtual Chassis Components.

6. Confirm that the new member switch is now included in the Virtual Chassis configuration on switches with a front-panel LCD by checking the display for the member ID. It should show a member ID in the range from 0 through 9.

   If you are using a switch that does not have an LCD interface, enter the`show virtual-chassis` command and view the output to confirm the switch is part of the Virtual Chassis configuration.

7. Cable the other VCP on the replacement member switch into the Virtual Chassis based on how you planned to interconnect the switch in Step 5 of this procedure.

   CAUTION: If you immediately cable both VCPs on the new switch into the existing Virtual Chassis at the same time, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. Network traffic to this switch is dropped during the downtime.

   The member switch will return to the normal operational state with no user intervention, and normal operation of the Virtual Chassis will resume after this downtime.

8. IF you need to update the new member switch's current member ID to the member ID of the switch that was removed from the Virtual Chassis configuration:

- In a nonprovisioned Virtual Chassis, issue the `request virtual-chassis renumber` command on the primary member switch.

- In a preprovisioned Virtual Chassis, on the primary member switch, reconfigure the member information for the new member switch using the `[edit virtual-chassis] member` configuration statement.

  To use the same member ID as the member that was replaced, associate the new switch's serial number (on the back of the switch) with the replaced member ID, as follows:

  ```
  [edit virtual-chassis]
  user@switch# set member replaced-member-ID serial-number new-member-serial-number
  ```

  NOTE: You can alternatively use the `replace` configuration editing command to substitute the serial number of the replacement member switch for the replaced member's serial number in the existing configuration item for the replaced member.

  To configure the new member switch with a different member ID, associate the new switch's serial number with the desired member ID and then delete the configuration item for the replaced member switch, as follows:

  ```
  [edit virtual-chassis]
  user@switch# set member new-member-ID serial-number new-member-serial-number
  user@switch# delete member replaced-member-ID
  ```

Replace a Member Switch With a Different Type of Switch That Changes the Virtual Chassis to Mixed Mode

If you want to replace a member switch with a different type of switch that changes the Virtual Chassis from a non-mixed to a mixed Virtual Chassis:

1. Remove the member switch as described in "Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch" on page 123.
2. Use the configuration procedure for *adding a new switch to an existing Virtual Chassis* based on the type of switch you are adding. (See the list of related documentation at the bottom of this page.)

See Understanding Mixed EX Series and QFX Series Virtual Chassis for the combinations of switches that comprise a mixed Virtual Chassis.

RELATED DOCUMENTATION

Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis

Adding an EX4600 Switch to a Mixed or Non-mixed Virtual Chassis

Adding a New Switch to an Existing EX4650 or QFX Series Virtual Chassis

Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis (CLI Procedure)

Adding an EX4200 Switch to a Preprovisioned EX4500 Virtual Chassis or a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis (CLI Procedure)

Adding an EX4500 Switch to a Preprovisioned EX4200 Virtual Chassis (CLI Procedure)

Adding an EX4500 Switch to a Nonprovisioned EX4200 Virtual Chassis (CLI Procedure)

Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis (CLI Procedure)

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port

Configuring Primary Role of a Virtual Chassis

IN THIS SECTION

NOTE: This topic applies to all QFX Virtual Chassis and all EX Series Virtual Chassis except EX8200 Virtual Chassis.

A Virtual Chassis configuration has two Routing Engines—one is the switch in the primary Routing Engine role and the other is the switch in the backup Routing Engine role. The remaining members operate in the linecard role. You can designate the role (primary, backup, or linecard) that a member switch performs within any Virtual Chassis whether or not you are using a preprovisioned configuration. For details on which switches in a mixed Virtual Chassis we recommend or require you to configure into the primary or backup Routing Engine role, see "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 54.

> **NOTE**: We recommend that you always use `commit synchronize` rather than `commit` to save configuration changes for a Virtual Chassis to ensure that the changes are saved on both Routing Engines.

Configuring Primary Role Using a Preprovisioned Configuration File

To configure primary role using a preprovisioned configuration:

1. Note the serial numbers of the switches that you want to function as the primary and backup Routing Engines.

   > **NOTE**: Serial number values are case-sensitive.

2. Power on only the switch that you want to function as the primary Routing Engine.
3. Edit the configuration to specify the preprovisioned configuration mode:

   ```
   [edit virtual-chassis]
   user@switch# set preprovisioned
   ```

4. Specify the serial numbers of the member switches that you want to function as primary and backup, specifying their role as `routing-engine`:

   ```
   [edit]
   user@switch# set virtual-chassis member 0 serial-number abc123 role routing-engine
   user@switch# set virtual-chassis member 1 serial-number def456 role routing-engine
   ```

   > **NOTE**: You cannot directly modify the primary-role priority value when you configure a member switch in a preprovisioned Virtual Chassis. The Virtual Chassis generates the primary-role priority values automatically. The values depend on the role you assign to the

member switches in the configuration file. The default primary-role priority on any switch is 128. When you configure the primary and backup Routing Engine members, the primary-role priority changes to 129 on those members. According to the primary-role election algorithm, the member that you powered on first gets priority and becomes the primary. See "Understanding How the Primary in a Virtual Chassis Is Elected" on page 62. You can configure only two members with the `routing-engine` role.

5. Specify the serial numbers of any other member switches you are including in the Virtual Chassis configuration. You can also explicitly configure their role as `line-card`.

Configuring Primary Role Using a Configuration File That Is Not Preprovisioned

To configure primary role of the Virtual Chassis through a configuration that is not preprovisioned:

1. Power on only the switch that you want to function as the primary Routing Engine.
2. Configure the highest possible primary-role priority value (`255`) for the member that you want to function as the primary Routing Engine:

```
[edit virtual-chassis]
user@switch# set member 0 mastership-priority 255
```

3. Configure the same primary-role priority value (continue to edit the Virtual Chassis configuration on the primary) for the member that you want to be the backup Routing Engine:

```
[edit virtual-chassis]
user@switch# set member 1 mastership-priority 255
```

NOTE: You should assign the same (highest) primary-role priority value to both the primary and backup Routing Engine members for reliable graceful Routing Engine switchover (GRES) operation. This and the other primary-role election considerations also help keep primary role from switching back and forth rapidly between the two members under failover conditions.

4. Use the default primary-role priority value (`128`) for the remaining member switches or configure the primary-role priority to a value that is lower than the value specified for members acting in the primary and backup Routing Engine roles.

RELATED DOCUMENTATION

Configuring the Timer for the Backup Member to Start Using Its Own MAC Address as Primary of a Virtual Chassis

When a backup member takes control of a Virtual Chassis because of a reset or other temporary failure, the backup member uses the MAC address of the old primary switch as the system MAC base address. This process helps ensure a smooth transition of the primary role with no disruption to network connectivity.

The MAC persistence timer is used in situations in which the primary switch is no longer a member of the Virtual Chassis because it has been physically disconnected or removed. If the old primary switch does not rejoin the Virtual Chassis before the timer elapses, the new primary switch starts using its own MAC address as the system's MAC base address. For information regarding how the system MAC base address is used to assign MAC addresses to ports in a Virtual Chassis, see "Understanding MAC Address Assignment on a Virtual Chassis" on page 76.

The default timer value is 10 minutes. The maximum timer value is 60 minutes.

You can disable the MAC persistence timer starting in Junos OS Release 12.1R3. When the MAC persistence timer is disabled, the MAC address of the old primary switch is used as the system MAC base address; no MAC address changes occur within the Virtual Chassis even when the old primary switch is no longer a member of the Virtual Chassis because it has been physically disconnected or removed.

To configure or modify the MAC persistence timer:

```
[edit virtual-chassis]
user@switch# set mac-persistence-timer minutes
```

To disable the MAC persistence timer:

```
[edit virtual-chassis]
user@switch# set mac-persistence-timer disable
```

RELATED DOCUMENTATION

Configuring an EX4650 or a QFX Series Virtual Chassis | 100
Understanding Virtual Chassis Components | 35

Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port

IN THIS SECTION

Use the procedures described in this topic to set up Virtual Chassis ports (VCPs) to connect two switches together in an EX Series or a QFX Series Virtual Chassis.

Switches that can be members of a Virtual Chassis might have:

- Dedicated VCPs—Ports you can use only as VCPs.

- Default-configured VCPs—Ports that are already configured into VCPs with the default factory configuration. On some switches these ports can alternatively be converted back into and used as uplink or network ports.

- Uplink or network ports that are also supported as VCPs—Ports you can configure into VCPs and convert back into uplink or network ports as needed.

If you don’t have or are not using dedicated or default-configured VCPs, you need to convert supported ports into VCPs to interconnect Virtual Chassis members. When a switch has uplink or network ports that can be converted into VCPs, you can use some ports as VCPs and others as network ports or uplinks to other devices in trunk mode. When you set a port as a VCP, you can’t use it for any other purpose.

NOTE: You don't use the procedures in this topic to set the VCPs on EX4400 switches. EX4400 switches have ports that operate as VCPs by default, and no other ports on the switch can be used as VCPs. If you change the operational mode of the default VCPs to network port mode to use them as network ports instead of as VCPs, to subsequently use the switch in a Virtual Chassis, you must disable network port mode to return those ports to their default VCP mode. To change the default VCPs to network port mode, use the `request virtual-chassis mode network-port <reboot>` command. To disable network port mode to convert the ports back into VCPs, use the `request virtual-chassis mode network-port disable <reboot>` command. (When you enable or disable network port mode, you must reboot the switch for the change to take effect.)

You can set ports as VCPs on a standalone switch before interconnecting any links into a Virtual Chassis, or set them after interconnecting one link on the switch into an existing Virtual Chassis. Either way, after the VCP is connected into the Virtual Chassis, the primary switch uses the link to detect the switch and complete the process of adding it as a member.

NOTE: Most Virtual Chassis also support *autoprovisioning*, which means that under certain conditions, when you interconnect a member switch into an existing Virtual Chassis, ports that are supported as VCPs will convert automatically into VCPs when you cable the link. This is an easy way to add member switches to a Virtual Chassis without needing to explicitly configure VCPs, but it only works if the ports on *both* sides of the link are *not already configured as VCPs*. If you want to use autoprovisioning, you might need to delete VCP settings (whether default-configured or those you previously set explicitly) on either or both sides of the links you are using to interconnect the new member switch. See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45 for details.

See "Virtual Chassis Port Options" on page 41 for a list of supported VCP ports on each type of switch. For complete details about where a switch has dedicated VCPs, default-configured VCPs, or ports that can be configured as VCPs, and what transceivers and cables are supported to use for VCP connections on that switch, see the hardware documentation for the switch. In general, even if a port is supported as a VCP, you can't use it as a VCP if it's channelized.

When to Configure VCPs

You typically configure a port as a Virtual Chassis Port (VCP) for one of the following reasons:

- You are configuring a Virtual Chassis composed of switches that support Virtual Chassis but do not have default-configured VCPs or dedicated VCPs.
- You are using default-configured VCPs or dedicated VCPs to interconnect members in a Virtual Chassis, and want to add redundant VCP links between members using additional ports that can be configured into VCPs.
- You want to interconnect two switches into a Virtual Chassis that have dedicated VCPs but are located in different wiring closets or sites, and the switches are farther apart then the maximum length of a dedicated VCP cable.
- You previously changed a default-configured VCP to use it as a network or uplink port, and now you want to use it as a VCP again.

We recommend that you have two uplink VCP connections within each wiring closet for redundancy. VCPs automatically bundle into a Link Aggregation Group (LAG) when two or more ports operating at the same speed are configured into VCPs between the same two member switches. See "Understanding Virtual Chassis Port Link Aggregation" on page 68 for details.

CAUTION: If you set a port as a VCP to create a redundant link with a dedicated VCP connection on EX4200, EX4500, or EX4550 switches, to avoid traffic looping within the Virtual Chassis, you must reboot the Virtual Chassis after configuring the port conversion and cabling the VCP link. See "Troubleshooting an EX Series Virtual Chassis" on page 162 for more information.

Prepare Virtual Chassis Member Switches Before Converting VCPs

Before converting a port into a VCP and interconnecting the switch into a Virtual Chassis:

1. Verify which ports can be used as VCPs in your particular configuration. See "Virtual Chassis Port Options" on page 41 for a summary of the VCP options on switches that support Virtual Chassis, and the hardware documentation for each type of switch for complete details about the ports and installed transceivers that can be used as VCPs.

2. If you are configuring an uplink module port as a VCP, if needed, install the uplink module in the member switches that you want to interconnect.
3. Log into the switch that is or will be the primary of the Virtual Chassis.

> **NOTE**: Do not power on the other switches at this point.

4. (EX Series switches only) Run EZSetup on the switch that you are configuring to be the primary member switch. Make sure the hostname and other identification, time zone, and network properties are set up on the primary. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* for details. The parameters you specify for the primary apply to the entire Virtual Chassis, including all the member switches that you interconnect later.
5. If you want to configure and manage the Virtual Chassis remotely, specify the VME global management interface. You can configure the VME global management interface when you are setting up the primary or you can do it after completing the other configuration steps for the Virtual Chassis. See *Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)*.
6. Configure primary role of the Virtual Chassis using either a nonprovisioned or preprovisioned configuration. See "Configuring Primary Role of a Virtual Chassis" on page 129 for details.

> **NOTE**: A Virtual Chassis has two Routing Engines, one in the primary role and the other in the backup role. Therefore, we recommend that you always use `commit synchronize` rather than simply `commit` to save configuration changes made for a Virtual Chassis. This ensures that the configuration changes are saved in both Routing Engines.

Before you begin to interconnect new Virtual Chassis members across long distances, such as between wiring closets:

- Prepare the existing Virtual Chassis for interconnecting with a potential member switch that is beyond the reach of a dedicated Virtual Chassis cable by setting at least one uplink VCP on an existing member of the Virtual Chassis.
- Prepare the potential member switch for interconnecting with the existing Virtual Chassis by setting at least one uplink VCP on the standalone switch.

Set Uplink Ports to Form VCP Links Between Member Switches in a Virtual Chassis

From the Virtual Chassis, you can set an uplink port on the local member or on a specified member as a VCP.

> **NOTE**: On EX4200 switches, if you use an SFP+ uplink module, you must configure all member switches to support either 1-gigabit SFP transceivers or 10-gigabit SFP+ transceivers. See Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module .

To set the uplink ports for the local member switch (for example, member 0) and for a different member switch (for example, member 1) to function as VCPs:

1. Set one uplink port of member 0 as a VCP. You don't need to specify the `member` *`member-id`* option, because the command applies by default on the member where it is executed. You can alternatively include the `local` option if you want to make sure the command applies only to that port locally on the switch where you're running the command.

   ```
   user@switch> request virtual-chassis vc-port set pic-slot 1 port 0
   ```

2. Set one uplink port of member 1 as a VCP. This step includes the member *member-id* option because it acts on a different member switch than the local member switch.

   ```
   user@switch> request virtual-chassis vc-port set pic-slot 1 port 0 member 1
   ```

   > **NOTE**: You can also connect to a member switch individually using the *request session member* command and set a VCP locally on that member. (You don't specify the `member` option in that case.)

Set an Uplink Port as a VCP on a Standalone Switch

You can set an uplink VCP on a standalone switch before interconnecting the link into an existing Virtual Chassis. You must set the port as a VCP for the Virtual Chassis primary to detect the switch and complete the process of adding it as a member.

To set one uplink VCP on the potential member, which is currently operating as a standalone switch:

1. Power on the standalone switch.
2. Set one uplink port as a VCP. You do not need to specify the `member` *`member-id`* option, because the command applies by default on the member where it is executed.

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port 0
```

   NOTE: If you do specify the `member` *`member-id`* option, use member ID `0`. Because the switch is not yet interconnected with the other members of the Virtual Chassis, its current member ID is `0`. Its member ID will change when it is interconnected with the Virtual Chassis. It does not impact the functioning of the uplink VCP that its VCP is set with `0` as the member ID. The VCP has significance only on the local switch.

3. After you have set the uplink VCP on the standalone switch, physically interconnect its uplink port with a VCP uplink port on one of the members in the existing Virtual Chassis.

   The new member switch reboots and joins the now expanded Virtual Chassis with a different member ID.

   NOTE: The setting for the new member switch's uplink VCP remains intact and is not affected by the change of member ID.

4. Repeat the steps above to interconnect the new switch to another member of the Virtual Chassis or to create redundant VCP links to the same member.

Remove a VCP Setting on an Uplink or Network Port

You might want to convert a VCP back into an uplink or network port if:

- You want to use a default-configured VCP on a switch as a network or uplink port instead of as a VCP.
- You want to add a new member switch or a new VCP link in an existing Virtual Chassis using autoprovisioning, where the VCP links form automatically when you cable them *only* if the ports on both sides of the link are not already set as VCPs. (See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45 for details.)
- You remove a switch from a Virtual Chassis and want to use it as a standalone switch again.

  NOTE: In this case, whenever possible we recommend that you revert the switch back to its default factory configuration to guarantee the smoothest transition back to standalone operation. See *Removing or Replacing a Member Switch of a Virtual Chassis Configuration*.

To remove the VCP setting on a port:

```
user@switch> request virtual-chassis vc-port delete pic-slot slot-number port port-number
```

You usually do this on the switch with the VCP itself, so you don't need to include the `member member-id` option because the command applies by default on the member where you run it.

RELATED DOCUMENTATION

Configuring an EX3300 Virtual Chassis (CLI Procedure)

Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)

Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)

Configuring an EX2300, EX3400, EX4100, EX4100-F, EX4300, or EX4400 Virtual Chassis | **80**

Configuring an EX4650 or a QFX Series Virtual Chassis | **100**

Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets

Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Disabling Split and Merge in a Virtual Chassis

The split and merge feature is enabled by default on all EX Series switches and QFX Series devices that are connected into a Virtual Chassis.

If a Virtual Chassis splits:

- Each part detects the members they can still reach and elects a new primary for those members to form a smaller Virtual Chassis.
- Each part determines if it will remain the active new configuration of the original Virtual Chassis. The feature's rules ensure only one part remains active after a split.

See "Understanding Split and Merge in a Virtual Chassis" on page 71 for complete details on what happens during a Virtual Chassis split or merge.

You can disable the split and merge feature in a Virtual Chassis by configuring the `no-split-detection` option on the Virtual Chassis. In this case, if the Virtual Chassis splits, both parts of the split Virtual Chassis configuration usually remain active.

However, with split detection disabled, be aware that the two resulting Virtual Chassis might not form as you might expect. For example, in a preprovisioned Virtual Chassis:

- If both of the Routing Engines end up in the same Virtual Chassis configuration after a split, the other part of the split Virtual Chassis configuration can't form and remains inactive (because it has no members that can take on the primary or backup Routing Engine).
- If the Routing Engines end up in different parts of the split Virtual Chassis configuration and the rest of the member switches are configured in linecard roles, then the resulting Virtual Chassis parts might not be able to select a backup Routing Engine.

BEST PRACTICE: We very strongly recommend that you:

- Keep the default setting with the split and merge feature enabled for any Virtual Chassis with more than two members.

  This provides the most predictable results when a split happens, and enables the split and merge feature to perform a merge after the issues that caused the split are resolved.

- Disable split and merge on a Virtual Chassis with only two member switches by setting the `no-split-detection` option.

We have found that a two-member switch Virtual Chassis with split and merge disabled re-forms more quickly with fewer complications when a split occurs.

If you expand a two-member Virtual Chassis to include more members, delete the `no-split-detection` setting to re-enable split and merge feature again.

To disable the split and merge feature in a Virtual Chassis:

```
[edit]
user@switch# set virtual-chassis no-split-detection
```

To remove this configuration item to restore the default configuration with the split and merge feature again—we very strongly recommend that you do this when a two-member Virtual Chassis is expanded to have more members:

```
[edit]
user@switch# delete virtual-chassis no-split-detection
```

RELATED DOCUMENTATION

Understanding Split and Merge in a Virtual Chassis | 71

Configuring Automatic Software Update on Virtual Chassis Member Switches

The automatic software update feature allows you to automatically update the software version on prospective member switches as they are added so that they can join the Virtual Chassis.

NOTE: The version of Junos OS running on the Virtual Chassis must be compatible with the software running on the prospective member switch for an automatic software update to occur. For information on Junos OS compatibility and other automatic software update restrictions, see "Understanding Automatic Software Update on Virtual Chassis Member Switches" on page 74.

Before you begin, ensure that you know the name or the URL of the software package to be used by the automatic software update feature.

To configure the automatic software update feature for an EX Series or QFX Series Virtual Chassis that is not a mixed-mode Virtual Chassis:

```
[edit]
user@switch# set virtual-chassis auto-sw-update package-name package-name
```

NOTE: An EX4300 Virtual Chassis with a combination of EX4300 multigigabit model switches (EX4300-48MP) and other EX4300 switches must be configured in mixed mode, and the automatic software update feature is not supported in this case.

NOTE: A QFX5110 Virtual Chassis is considered to be a non-mixed Virtual Chassis. Because both QFX5110 and QFX5100 switches in a QFX5110 Virtual Chassis run the same software image, you can use the `auto-sw-update` command with one software package name for all members.

CAUTION: A QFX5100 switch running a Junos OS software image with "-qfx-5-" in the package filename *must* first be upgraded to a Junos OS software image with "-qfx-5e-" in the package filename before it can be added to a QFX5110 Virtual Chassis or VCF. The automatic software update process cannot update a switch from a "-qfx-5-" image to a "-qfx-5e-" image. See *Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric*.

After a QFX5100 switch is installed with a "-qfx-5e-" Junos OS software image, the automatic software update process can successfully update the switch automatically with a different version of a "-qfx-5e-" Junos OS image to match the other members in the Virtual Chassis or VCF.

To configure the automatic software update feature on a mixed Virtual Chassis containing at least one EX4200 switch and at least one EX4500 or EX4550 switch, use the `ex-4200` option when you are specifying a path to a package for the EX4200 switches and the `ex-4500` option when you are specifying a path to a package for the EX4500 or EX4550 switches, as follows:

```
[edit]
user@switch# set virtual-chassis auto-sw-update ex-4200 package-name package-name
user@switch# set virtual-chassis auto-sw-update ex-4500 package-name package-name
```

You do not need to specify the `ex4500` option when configuring automatic software update for a Virtual Chassis with only a combination of EX4500 and EX4550 switches, which does not need to run in mixed mode. Specifying only a Junos OS package name for an EX4500 switch updates the software for both EX4500 and EX4550 switches.

To configure the automatic software update feature on a mixed QFX5100 Virtual Chassis composed of QFX5100 switches and at least one other supported type of device (QFX3500, QFX3600, or EX4300 switches), you must specify a software package name for each type or family of device in the mixed Virtual Chassis:

- Specify the `qfx-5` option with the path to the Junos OS package for QFX5100 switches
- Specify the `qfx-3` option with the path to the Junos OS package for QFX3600 and QFX3500 switches
- Specify the `ex4300` option with the path to the Junos OS package for EX4300 switches

as follows:

```
[edit]
user@device# set virtual-chassis auto-sw-update qfx-5 package-name package-name
user@device# set virtual-chassis auto-sw-update qfx-3 package-name package-name
user@device# set virtual-chassis auto-sw-update ex-4300 package-name package-name
```

You can similarly specify different package names using the `ex-4600` and `ex-4300` options with the `auto-sw-update` statement to configure the automatic software update feature for a mixed EX4600 Virtual Chassis that contains EX4600 and EX4300 member switches.

If the software package is located on a local directory on the switch, use the following format for *package-name*:

/pathname/package-name

If the software package is to be downloaded and installed from a remote location, use one of the following formats:

`ftp://`*hostname/pathname/package-name*

`ftp://`*username*:*prompt*`@ftp.`*hostname*`.net/`*package-name*

`http://`*hostname/pathname/package-name*

RELATED DOCUMENTATION

Understanding Automatic Software Update on Virtual Chassis Member Switches | 74

Understanding Mixed EX Series and QFX Series Virtual Chassis | 54

Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge

Every Virtual Chassis has a unique ID that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a Virtual Chassis ID using the **set virtual-chassis id** command. When two Virtual Chassis configurations attempt to merge, the Virtual Chassis ID that you assigned takes precedence over the automatically assigned Virtual Chassis IDs and becomes the ID for the newly merged Virtual Chassis configuration.

To configure the Virtual Chassis ID:

```
[edit]
user@switch# set virtual-chassis id                              id
```

RELATED DOCUMENTATION

Understanding Split and Merge in a Virtual Chassis | 71

Configuring Graceful Routing Engine Switchover in a Virtual Chassis

In a Virtual Chassis, one member switch is assigned the primary role and has the primary Routing Engine. Another member switch is assigned the backup role and has the backup Routing Engine. Graceful Routing Engine switchover (GRES) enables the primary and backup Routing Engines in a Virtual Chassis configuration to switch from the primary to backup without interruption to packet forwarding as a hitless failover solution. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the primary Routing Engine to preserve kernel state information and the forwarding state.

To set up the Virtual Chassis configuration to use graceful Routing Engine switchover (GRES):

1. Set up a minimum of two switches in a Virtual Chassis configuration with primary-role priority of 255:

```
[edit]
user@switch# set virtual-chassis member 0 mastership-priority 255

[edit]
user@switch# set virtual-chassis member 1 mastership-priority 255
```

2. Set up graceful Routing Engine switchover:

```
[edit]
user@switch# set chassis redundancy graceful-switchover
```

Commit the configuration.

> **NOTE**: We recommend that you use the `commit synchronize` command to save any configuration changes that you make to a multimember Virtual Chassis.

RELATED DOCUMENTATION

Example: Configuring an EX4200 Virtual Chassis with a Primary and Backup in a Single Wiring Closet

High Availability Features for EX Series Switches Overview

Understanding EX Series Virtual Chassis | 9

Understanding QFX Series Virtual Chassis | 25

3

CHAPTER

Virtual Chassis Routine Monitoring and Troubleshooting

Command Forwarding Usage with EX Series and QFX Series Virtual Chassis

Some CLI commands can be run either on all members or on a specific member of a Virtual Chassis configuration. This functionality is referred to as command forwarding.

You can always specify that these commands be applied to all member switches in the Virtual Chassis by using the `all-members` option, or to a specific member switch by using the `member-`*`member-id`* option. If neither option is specified, the default command forwarding behavior, which varies by command, is used. See the **Default** row in Table 8 on page 148 to learn the command forwarding behavior for a specific command.

For example, to collect information about a particular member switch prior to contacting Juniper Networks Technical Assistance Center (JTAC), use the `request support information member` *`member-id`* command to gather data for the specified member switch. If you want to gather this data for all member switches in the Virtual Chassis, you can enter the `request support information` command, which by default uses the `all-members` option, or the `request support information all-members` command.

Table 8 on page 148 provides a list of commands that can be run either on all members of the Virtual Chassis configuration or on a specific member switch.

Table 8: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration

Command	Purpose	`all-members`	`member-`*`member-id`*	Default
`request support information`	Use this command when you contact JTAC about your component problem. This command is the equivalent of using the following CLI commands: • `show version` • `show chassis firmware` • `show chassis hardware` • `show chassis environment` • `show interfaces extensive` (for each configured interface) • `show configuration` (excluding any SECRET-DATA) • `show system virtual-memory`	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`

Table 8: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration *(Continued)*

Command	Purpose	`all-members`	`member-`*`member-id`*	Default
`request system partition hard-disk`	Set up the hard disk for partitioning. After this command is issued, the hard disk is partitioned the next time the system is rebooted. When the hard disk is partitioned, the contents of /**altroot** and /**altconfig** are saved and restored. All other data on the hard disk is at risk of being lost.	Partitions the hard disk on all members of the Virtual Chassis configuration.	Partitions the hard disk on the specified member switch.	`all-members`
`request system reboot`	Reboot Junos OS for EX Series or QFX Series switches after a software upgrade and occasionally to recover from an error condition.	Reboots all members of the Virtual Chassis configuration.	Reboots the specified member switch.	`all-members`
`request system snapshot`	Back up the currently running and active file system.	Backs up the file systems on all members of the Virtual Chassis configuration.	Backs up the file system on the specified member switch.	`all-members`

Table 8: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration *(Continued)*

Command	Purpose	`all-members`	`member-`*`member-id`*	Default
`request system storage cleanup`	Free storage space on the switch by rotating log files and proposing a list of files for deletion. User input is required for file deletion.	Runs cleanup on all members of the Virtual Chassis configuration.	Runs cleanup on the specified member switch.	`all-members`
`show log user`	Display users who are viewing the system log.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	primary switch only
`show system alarms`	Display active system alarms.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system audit`	Display the state and checksum values for file systems.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`

Table 8: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration *(Continued)*

Command	Purpose	`all-members`	`member-*member-id*`	Default
`show system boot-messages`	Display initial messages generated by the system kernel upon startup. These messages are the contents of **/var/run/dmesg.boot**.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system buffers`	Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on).	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`

Table 8: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration *(Continued)*

Command	Purpose	`all-members`	`member-`*`member-id`*	Default
`show system connections`	Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and which connections are currently in progress.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system core-dumps`	Display a core file generated by an internal Junos OS process.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system directory-usage`	Display directory usage information.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	primary switch only
`show system processes`	Display information about software processes that are running on the switch and that have controlling terminals.	Displays information for all members of the Virtual Chassis configuration.		`all-members`

Table 8: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration *(Continued)*

Command	Purpose	`all-members`	`member-member-id`	Default
`show system reboot`	Display pending system reboots or halts.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system snapshot`	Display information about the backup software that is located in the **/ altroot** and **/ altconfig** file systems. To back up software, use the request system snapshot command.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system software`	Display the Junos OS extensions loaded on your switch.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system statistics`	Display systemwide protocol-related statistics.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`

Table 8: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration *(Continued)*

Command	Purpose	`all-members`	`member-`*`member-id`*	Default
`show system storage`	Display statistics about the amount of free disk space in the switch's file systems.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system uptime`	Display the current time and information about how long the switch, the switch software, and any existing protocols have been running	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`
`show system users`	Show all users who are currently logged in.	Shows all users who are currently logged in to any members of the Virtual Chassis configuration.	Shows all users who are currently logged in to the specified member switch.	`all-members`
`show system virtual-memory`	Display the usage of Junos OS kernel memory, listed first by size of allocation and then by type of usage. Use `show system virtual-memory` for troubleshooting with JTAC.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.	`all-members`

RELATED DOCUMENTATION

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis
Virtual Chassis Overview for Switches | 2
Understanding Virtual Chassis Components | 35

Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member

IN THIS SECTION

Purpose

You can designate the role that a member performs within a Virtual Chassis or you can allow the role to be assigned by default. You can designate the member ID that is assigned to a specific switch by creating a permanent association between the switch's serial number and a member ID, using a preprovisioned configuration. Or you can let the member ID be assigned by the primary, based on the sequence in which the member switch is powered on and on which member IDs are currently available.

The role and member ID of the member switch are displayed on the front-panel LCD (for switches that have an LCD) or in the output from the `show virtual-chassis` CLI command.

Each member switch can be cabled to one or two other member switches, using either the dedicated Virtual Chassis ports (VCPs) on the rear panel, an uplink port that has been configured as a VCP, or an optical port that has been configured as a VCP. The members that are cabled together are considered neighbor members.

Action

To display the role and member ID assignments using the CLI:

```
user@switch>                      show virtual-chassis

Virtual Chassis ID: 0000.e255.00e0

                                       Mastership            Neighbor List
Member ID  Status   Serial No    Model     Priority   Role        ID, Interface

0 (FPC 0)  Prsnt    abc123       ex4200-48p    255   Master*      1 vcp-0
                                                                  2 vcp-1

1 (FPC 1)  Prsnt    def456       ex4200-24t    255   Backup       2 vcp-0
                                                                  0 vcp-1

2 (FPC 2)  Prsnt    abd231       ex4200-24p    128   Linecard     0 vcp-0
                                                                  1 vcp-1
```

Meaning

This output verifies that three EX4200 switches have been interconnected as a Virtual Chassis configuration through their dedicated VCPs to create an EX4200 Virtual Chassis. The display shows which of the VCPs is connected to which neighbor. The first port (**vcp-0**) of member **0** is connected to member **1** and the second port of member **0** (**vcp-1**) is connected to member **2**. The FPC slots for the switches are the same as the member IDs.

The **Mastership Priority** values indicate that the primary and backup members have been explicitly configured, because they are not using the default value (**128**).

> **NOTE**: This example uses output from an EX4200 Virtual Chassis. The output, with the exception of the **Model** column, would be identical on all other Virtual Chassis.

RELATED DOCUMENTATION

Configuring Primary Role of a Virtual Chassis | 129
Configuring an EX4650 or a QFX Series Virtual Chassis | 100
Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Verifying That Virtual Chassis Ports Are Operational

IN THIS SECTION

Purpose

Display the status of Virtual Chassis ports (VCPs) in a Virtual Chassis or Virtual Chassis Fabric (VCF).

NOTE: VCPs are not displayed when you issue the `show interfaces` command.

Action

Display the VCPs:

```
user@switch> show virtual-chassis vc-port all-members


fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
```

```
or                          ID               (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated        1     Up         32000        1   vcp-0
vcp-1       Dedicated        2     Up         32000        1   vcp-1
1/0         Configured       3     Up         1000         2   vcp-255/1/0
1/1         Configured       3     Up         1000         2   vcp-255/1/1
1/2         Configured       4     Up         1000         4   vcp-255/0/20
1/3         Configured       4     Up         1000         4   vcp-255/0/21

fpc1:
--------------------------------------------------------------------------
Interface   Type            Trunk  Status      Speed        Neighbor
or                          ID               (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated        1     Up         32000        0   vcp-0
vcp-1       Dedicated        2     Up         32000        0   vcp-1
1/0         Configured       3     Up         10000        3   vcp-255/1/0
1/1         Configured       3     Up         10000        3   vcp-255/1/1

fpc2:
--------------------------------------------------------------------------
Interface   Type            Trunk  Status      Speed        Neighbor
or                          ID               (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated        1     Up         32000        3   vcp-0
vcp-1       Dedicated        2     Up         32000        3   vcp-1
1/0         Configured       3     Up         1000         0   vcp-255/1/0
1/1         Configured       3     Up         1000         0   vcp-255/1/1
1/2                         -1     Down       1000
1/3                         -1     Down       1000

fpc3:
--------------------------------------------------------------------------
Interface   Type            Trunk  Status      Speed        Neighbor
or                          ID               (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated        1     Up         32000        2   vcp-0
vcp-1       Dedicated        2     Up         32000        2   vcp-1
1/0         Configured       3     Up         10000        1   vcp-255/1/0
1/1         Configured       3     Up         10000        1   vcp-255/1/1

fpc4:
--------------------------------------------------------------------------
```

```
Interface    Type             Trunk  Status       Speed        Neighbor
or                             ID                 (mbps)       ID   Interface
PIC / Port
vcp-0        Dedicated          1    Down         32000
vcp-1        Dedicated          2    Down         32000
0/20         Configured         3    Up           1000         0    vcp-255/1/2
0/21         Configured         3    Up           1000         0    vcp-255/1/3
```

Meaning

The dedicated VCPs are displayed as **vcp-0** and **vcp-1**. The uplink interfaces that have been set as uplink VCPs are displayed as **1/0**, **1/1**, **1/2**, and **1/3**. The network interfaces that have been set as VCPs are displayed as **0/20** and **0/21**. The neighbor interface names of uplink and network VCPs are of the form **vcp-255/*pic*/*port***—for example, **vcp-255/1/0**. In that name, **vcp-255** indicates that the interface is a VCP, **1** is the uplink PIC number, and **0** is the port number. The **fpc** number is the same as the member ID. The trunk ID is a positive number ID assigned to the link aggregation group (LAG) formed by the Virtual Chassis. If no LAG is formed, the value is **–1**.

NOTE: This example uses output from an EX4200 Virtual Chassis. The output is similar on all other types of Virtual Chassis or for a VCF.

RELATED DOCUMENTATION

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Configuring an EX3300 Virtual Chassis (CLI Procedure)

Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)

Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)

Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)

Verifying That Graceful Routing Engine Switchover Is Working in the Virtual Chassis

IN THIS SECTION

Purpose

Verify that a Graceful Routing Engine switchover (GRES) between two member switches acting as the primary and backup routing engines in a Virtual Chassis has occurred.

Action

On the primary switch, verify the member ID of the backup Routing Engine:

```
{master:0}
user@switch> show virtual-chassis
Virtual Chassis ID: 5efa.4b7a.aae6
                                          Mastership          Neighbor List
Member ID  Status   Serial No    Model    priority    Role      ID  Interface
0 (FPC 0)  Prsnt    BM0208105281 ex4200-24t      255   Master*    1  vcp-0
1 (FPC 1)  Prsnt    BP0208192350 ex4200-48t      255   Backup     0  vcp-0

Member ID for next new member: 2 (FPC 2)
```

1. Connect to the backup Routing Engine:

```
{master:0}
user@switch> request session member 1

{backup:1}
user@switch>
```

2. Verify that the backup Routing Engine is ready for switchover on member ID 1:

```
{backup:1}
user@switch> show system switchover

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

3. Switch the current backup Routing Engine to primary Routing Engine:

> **NOTE**: You must wait a minimum of two minutes between Routing Engine failovers for the Routing Engines to synchronize.

```
{backup:1}
user@switch> request chassis routing-engine master acquire
```

4. Verify that the primary and backup Routing Engines have switched roles:

> **NOTE**: Member ID **1** is now the primary and member ID **0** is now the backup.

```
{master:1}
user@switch> show virtual-chassis

Virtual Chassis ID: 5efa.4b7a.aae6
```

```
                                        Mastership           Neighbor List
Member ID  Status   Serial No    Model     priority    Role       ID  Interface
0 (FPC 0)  Prsnt    BM0208105281 ex4200-24t      255  Backup      1  vcp-0
1 (FPC 1)  Prsnt    BP0208192350 ex4200-48t      255  Master*     0  vcp-0

Member ID for next new member: 2 (FPC 2)
```

Meaning

With graceful Routing Engine switchover enabled, when you initiated a switchover from the backup Routing Engine, the backup Routing Engine became the primary and the primary Routing Engine became the backup.

RELATED DOCUMENTATION

Configuring Graceful Routing Engine Switchover in a Virtual Chassis

Troubleshooting an EX Series Virtual Chassis

IN THIS SECTION

This topic describes the following troubleshooting issues for a Virtual Chassis:

A Disconnected Member Switch's ID Is Not Available for Reassignment

IN THIS SECTION

Problem

Description

You disconnected a switch from the Virtual Chassis, but the disconnected switch's member ID is still displayed in the status output. You cannot reassign that member ID to another switch.

Solution

When you disconnect a member of a Virtual Chassis configuration, the primary retains the member ID and member configuration in its configuration database. Output from the `show virtual-chassis` command continues to display the member ID of the disconnected member with a status of `NotPrsnt`.

If want to permanently disconnect the member switch, you can free up the member ID by using the `request virtual-chassis recycle` command. This will also clear the status of that member.

Load Factory Default Does Not Commit on a Multimember Virtual Chassis

IN THIS SECTION

Problem

Description

The `load factory-default` command fails on a multimember Virtual Chassis.

Solution

The `load factory-default` command is not supported on a multimember Virtual Chassis configuration. For information on how to revert the switches in the Virtual Chassis to factory default settings, see *Reverting to the Default Factory Configuration for the EX Series Switch*.

The Member ID Persists When a Member Switch Is Disconnected From a Virtual Chassis

IN THIS SECTION

Problem

Description

Gigabit Ethernet interfaces retain their previous slot numbers when a member switch is disconnected from the Virtual Chassis.

Solution

If a switch had been previously connected as a member of a Virtual Chassis configuration, it retains the member ID that it was assigned as a member of that configuration even after it is disconnected and operating as a standalone switch. The interfaces that were configured while the switch was a member of the Virtual Chassis configuration retain the old member ID as the first digit of the interface name.

For example, if the switch was previously member 1, its interfaces are named `ge-1/0/0` and so on.

To change the switch's member ID, so that its member ID is 0, and to rename the switch's interfaces accordingly:

1. To change the member ID to 0:

```
user@switch> request virtual-chassis renumber member-id 1 new-member-id 0
```

2. To rename the interfaces to match the new member ID:

```
[edit virtual-chassis]
user@switch# replace pattern ge-1/ with ge-0/
```

A Member Switch Is Not Participating in a Mixed Virtual Chassis

IN THIS SECTION

Problem

Description

A member switch in a mixed Virtual Chassis is not participating in the Virtual Chassis. The `show virtual-chassis` output indicates the member switch status is `Inactive` or `NotPrsnt`.

This issue is most likely to occur immediately after you have cabled a mixed Virtual Chassis.

Solution

The Virtual Chassis mode on the switch might not be set to `mixed` mode. If the member switch is an EX4500 switch and is cabled into the Virtual Chassis through the dedicated Virtual Chassis port (VCP), the PIC mode might also be set to `Intraconnect` instead of `virtual-chassis`.

To verify the Virtual Chassis mode:

```
user@switch> show virtual-chassis mode

fpc0:
--------------------------------------------------------------------------
Mixed Mode: Enabled
fpc1:
--------------------------------------------------------------------------
Mixed Mode: Enabled
fpc2:
--------------------------------------------------------------------------
Mixed Mode: Enabled
fpc3:
--------------------------------------------------------------------------
Mixed Mode: Enabled
fpc4:
--------------------------------------------------------------------------
Mixed Mode: Disabled
fpc5:
--------------------------------------------------------------------------
Mixed Mode: Enabled
```

To change the Virtual Chassis mode on a member switch (in this case, member ID 4) to `mixed` mode:

```
user@switch> request virtual-chassis mode mixed member 4
```

(EX4500 switch only) To verify the PIC mode:

```
user@switch> show chassis pic-mode
fpc0:
--------------------------------------------------------------------------
    Pic Mode: Not-Applicable
fpc1:
--------------------------------------------------------------------------
    Pic Mode: Not-Applicable
fpc2:
--------------------------------------------------------------------------
    Pic Mode: Not-Applicable
fpc3:
--------------------------------------------------------------------------
```

```
    Pic Mode: Not-Applicable
fpc4:
--------------------------------------------------------------------------
    Pic Mode: PIC 3: Intraconnect
fpc5:
--------------------------------------------------------------------------
    Pic Mode: PIC 3: virtual-chassis
```

To change the PIC mode on an EX4500 switch to `virtual-chassis` mode (in this case, member ID 4):

```
user@switch> request chassis pic-mode virtual-chassis member 4
```

The member switch must be rebooted for the Virtual Chassis mode or PIC mode setting change to take effect. To reboot the member switch (in this case, member ID 4):

```
user@switch> request system reboot member 4
```

Unknown Traffic Looping Occurs After Configuring an Uplink Port as a Redundant VCP with a Dedicated VCP

IN THIS SECTION

Problem

Description

In a Virtual Chassis comprised of EX4200, EX4500, or EX4550 switches, you observe unrecoverable looping of unknown unicast or multicast traffic following the addition of a redundant VCP link between two member switches, when the two members are connected by a dedicated VCP link and the redundant link was created by converting uplink ports to VCPs.

This behavior can occur whether the redundant VCP link is created by setting the ports manually as VCPs or if the automatic VCP conversion feature is invoked and converts the ports into VCPs automatically.

Solution

Reboot the Virtual Chassis to properly detect the converted VCP as a redundant link with the dedicated VCP link.

After the conversion from a network port to a VCP, the egress filter table is not updated and the redundant VCP remains enabled for forwarding, which causes the looping behavior. The reboot process detects the converted port as a VCP and brings it up as disabled for forwarding.

As a result, we do not recommend connecting redundant converted uplink VCP ports between members already connected by dedicated VCPs on an active Virtual Chassis; instead, plan to add redundant uplink VCP connections during a maintenance window that can include a Virtual Chassis reboot cycle. This recommendation also applies when adding a new member to an existing active Virtual Chassis where you are adding redundant VCP links between the new member and one of its neighbors that mix dedicated VCPs and converted uplink VCPs.

RELATED DOCUMENTATION

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)

Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)

Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)

4

CHAPTER

Upgrading Software on a Virtual Chassis

Understanding Software Upgrades in a Virtual Chassis

IN THIS SECTION

This topic discusses software upgrades on EX Series and QFX Series *Virtual Chassis*, except EX8200 Virtual Chassis. For information on software upgrades on an EX8200 Virtual Chassis, see *Understanding Software Upgrades in an EX8200 Virtual Chassis*. For information on software upgrades on a Virtual Chassis Fabric (VCF), see *Understanding Software Upgrades in a Virtual Chassis Fabric*.

In a Virtual Chassis, each member switch must be running the same version of Juniper Networks Junos operating system (Junos OS) that supports Virtual Chassis. You can install a new Junos OS release on the entire Virtual Chassis or on individual members in the Virtual Chassis by using the same CLI command that you use to install Junos OS on standalone switches—the `request system software add` command.

In a mixed Virtual Chassis, the member switches must also be running the same version of Junos OS, but you might need to specify multiple Junos OS images when manually or automatically upgrading a mixed Virtual Chassis. For example, for the same Junos OS release, an EX4200 switch runs a different Junos OS image than an EX4500 or EX4550 switch, or an EX4300 switch runs a different Junos OS image than a QFX5100 switch in a QFX5100 Virtual Chassis. You can upgrade all member switches simultaneously by specifying a path to multiple Junos OS images in the same `request system software add` command.

NOTE: Some different switch platforms can run the same Junos OS image. For example, you do not need to use multiple Junos OS images when updating a Virtual Chassis with a combination of EX4500 and EX4550 switches, or a Virtual Chassis with a combination of QFX3500 and QFX3600 switches.

You can also use the following features to upgrade software on members of a Virtual Chassis:

Automatic Software Updates

You can use the automatic software update feature on a non-mixed or mixed Virtual Chassis to automatically update the Junos OS version on member switches as you add them to the Virtual Chassis. See "Understanding Automatic Software Update on Virtual Chassis Member Switches" on page 74 for more information.

If you are not configuring the automatic software update feature, we recommend that you update the new member switch to the version of Junos OS running on the Virtual Chassis before adding the member switch to the Virtual Chassis.

Nonstop Software Upgrade

You can also use nonstop software upgrade (NSSU) to upgrade Junos OS on all members of a Virtual Chassis for EX Series and QFX Series Virtual Chassis that support NSSU. NSSU provides an orderly upgrade of each member of the Virtual Chassis and takes advantage of *graceful Routing Engine switchover*, *nonstop active routing*, and link aggregation to minimize traffic disruption during the upgrade.

For more information about NSSU, see:

- "Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis" on page 178
- "Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade" on page 189

- (For legacy EX Series switches in a Virtual Chassis) *Upgrading Software Using Nonstop Software Upgrade on EX Series Virtual Chassis and Mixed Virtual Chassis (CLI Procedure)*

If NSSU is not supported to upgrade a Virtual Chassis from the currently-installed release to the release you need, refer to procedures like these for some guidance:

- Two-Member QFX Series Virtual Chassis Upgrade Procedure, a network configuration example on how to manually upgrade a two-member QFX Series Virtual Chassis when NSSU isn't available.

RELATED DOCUMENTATION

Understanding Virtual Chassis Components | 35

Configuring Automatic Software Update on Virtual Chassis Member Switches | 141

Installing Software on an EX Series Switch with a Virtual Chassis or Single Routing Engine (CLI Procedure)

Installing Software on a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Switches (CLI Procedure)

Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric

IN THIS SECTION

Use this procedure to upgrade a standalone QFX5100 switch running "QFX 5 Series" Junos OS software to a "QFX 5e Series" software image so the QFX5100 switch can join a QFX5110 Virtual Chassis or Virtual Chassis Fabric (VCF).

A QFX5110 Virtual Chassis or VCF can have a combination of QFX5110 and QFX5100 switches all running compatible Junos OS "QFX 5e Series" software. This procedure explains how you identify the compatible "QFX 5e Series" install package for QFX5100 switches, create a USB boot device as the installation media, and use the boot device to upgrade the software on a QFX5100 switch to the "QFX 5e Series" image. After installing the new software image, when you reboot the QFX5100, the switch is running a compatible software image and you can successfully add it into a QFX5110 Virtual Chassis or VCF.

NOTE: With releases prior to Junos OS Release 17.3R2 or 17.4R1, you must use this USB install method due to the differences in the boot structure and host OS software of the two types of devices.

Starting in Junos OS Releases 17.3R2 and 17.4R1, you are not required to use the USB install method to upgrade a QFX5100 switch from a “QFX 5 Series” to a “QFX 5e Series” image. Instead, you can install the “QFX 5e Series” package directly using the CLI command *request system software add*. See Installing a Standard Software Package on QFX5000 and EX4600 Switches for details on using the CLI install method.

If a QFX5100 switch is already running a “QFX 5e Series” image (software package filename contains the string “-qfx-5e-”), a QFX5110 Virtual Chassis or VCF can automatically update it to the right release when you add it to the Virtual Chassis or VCF. See Understanding Software Upgrades in a Virtual Chassis and *Understanding Software Upgrades in a Virtual Chassis Fabric*.

Identifying Compatible Software for QFX5100 Switches to Run in a QFX5110 Virtual Chassis or VCF

Standalone QFX5100 switches traditionally run Junos OS “QFX 5 Series” software, and the corresponding install media and installation software package file names include the string “-qfx-5-”. For example:

```
install-media-qfx-5-17.2R1.13.tgz
jinstall-host-qfx-5-17.2R1.13-signed.tgz
```

QFX5110 switches run Junos OS “QFX 5e Series” software and use a secure-boot method at startup, so the install media and software package filenames for QFX5110 switches include “-qfx-5e-” and “secure”. For example:

```
install-media-host-usb-qfx-5e-x86-64-17.2R1.13-secure.tgz
jinstall-host-qfx-5e-x86-64-17.2R1.13-secure-signed.tgz
```

QFX5100 switches can also run “5e Series” software, but in releases prior to Junos OS Release 17.3R2 or 17.4R1 they do not use the same secure boot method as QFX5110 switches, so the “QFX 5e Series” install media and software package filenames for QFX5100 switches include “-qfx-5e-” without the “secure” keyword. For example:

```
install-media-host-usb-qfx-5e-x86-64-17.3R1.6-signed.tgz
```

To add a QFX5100 switch to a QFX5110 Virtual Chassis or VCF, the QFX5100 must be running the same “QFX 5e Series” software version as the other switches in the Virtual Chassis or VCF. If the

QFX5100 switch has a "QFX 5 Series" image, you must first upgrade it manually to a "QFX 5e Series" image using a USB boot device that does not employ the secure-boot method.

To create the USB boot device (see "Creating a USB Boot Device for a QFX5100 Switch" on page 174), use the same install media filename *without* the "secure" keyword that matches what's running on the other QFX5110 Virtual Chassis or VCF members. For example:

If the Virtual Chassis or VCF is running the software from this install media package for QFX5110 switches (with the secure-boot method):

```
install-media-host-usb-qfx-5e-x86-64-17.3R1.6-secure-signed.tgz
```

Then the matching install media package for QFX5100 switches (without the secure-boot method) is:

```
install-media-host-usb-qfx-5e-x86-64-17.3R1.6-signed.tgz
```

After any QFX5100 switches are running a "QFX 5e Series" image, you can just use the same "jinstall-host-qfx-5e-" package file *with* the "secure" keyword to update the "QFX 5e Series" software running on all members of the Virtual Chassis or VCF, because when it starts up, the secure-boot install software determines whether or not to use the secure-boot method based on the type of switch on which it's running. The Virtual Chassis or VCF can also successfully update any "QFX 5e Series" member switches as needed in the same way with the automatic software update feature for adding or replacing members, or even during initial Virtual Chassis or VCF configuration.

NOTE: If you remove a QFX5100 switch from a QFX5110 Virtual Chassis or VCF and want to revert the QFX5100 switch to a "QFX 5 Series" software image to run as a standalone switch, you need to reinstall the image on the switch using a USB boot device with a "qfx-5-" install media file that does not use the secure-boot method.

Creating a USB Boot Device for a QFX5100 Switch

Use the following procedure to create a USB boot device with a Junos OS "QFX 5e Series" install media package (contains "-qfx-5e-" in the package filename) for a QFX5100 switch. You can then use the USB boot device to upgrade a QFX5100 switch to run that image.

NOTE: You can create the USB boot device on the switch you want to upgrade, on another Juniper Networks switch or router, or on any PC or laptop that supports Linux. The following steps describe creating the boot device from a Junos OS device, and might differ based on the device you use to create the boot device.

Before you begin, download the installation media file from https://www.juniper.net/customers/support/ to the device where you are creating the boot device. The install media filename should include the string "-qfx-5e-" for the same Junos OS release as the QFX5110 Virtual Chassis or VCF, but without a "secure" keyword in the filename, because you must initially upgrade QFX5100 switches without using the same secure-boot method employed by QFX5110 switches. (See "Identifying Compatible Software for QFX5100 Switches to Run in a QFX5110 Virtual Chassis or VCF" on page 173 for details.)

NOTE: The Junos OS software running on the QFX5110 members must be the "QFX 5e Series" image that uses the secure-boot method, so the install media and package filenames used on QFX5110 switches include the "secure" keyword. After initial USB installation of a "QFX 5e Series" software image on a QFX5100 switch, the same software image that supports secure boot runs on either switch model, and determines the appropriate boot method to use based on the switch on which it is running. As a result, for future updates, you can directly install the same image with the "secure" keyword on both QFX5110 and QFX5100 switches in your Virtual Chassis or VCF.

On a Junos OS device where you are creating the USB boot device:

1. Use FTP to copy the installation media file into the **/var/tmp** directory.
2. Insert the USB storage device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Switch to the root account using the `su` command:

```
% su
Password: password
```

NOTE: The password is the root password for the device on which you are creating the boot media. If you logged in to the device as root, you do not need to perform this step.

5. (Optional) Before copying the installation media file to the USB device, erase the boot sector of the USB device. In some cases, depending on how the USB device was formatted previously, this step can help avoid unexpected behavior during the USB boot process. For example, enter the following command:

```
root@device% dd if=/dev/zero of=/dev/da1 count=20
20+0 records in
20+0 records out
10240 bytes (10 kB) copied, 0.008281 seconds, 1.2 MB/s
```

6. Enter the following command to copy the installation media file (see "Identifying Compatible Software for QFX5100 Switches to Run in a QFX5110 Virtual Chassis or VCF" on page 173) to the USB device:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=1m
```

The device writes the installation media image to the USB device. For example:

```
root@device% dd if=install-media-host-qfx-5e-17.3R1.5-domestic.img of=/dev/da0 bs=1m
1399+0 records in
1399+0 records out
1466957824 bytes transferred in 394.081902 secs (3722469 bytes/sec)
```

7. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

8. Remove the USB storage device from the USB port.

You can now use the USB storage device to install the “QFX 5e Series” Junos OS software image on a QFX5100 switch.

Upgrading a QFX5100 Switch from Junos OS “QFX 5 Series” to “QFX 5e Series” Software Using a USB Boot Device

You must upgrade a standalone QFX5100 switch running “QFX 5 Series” software to “QFX 5e Series” software before the switch can join a QFX5110 Virtual Chassis or VCF. For this upgrade, you need to boot and install the software using a USB boot device.

Before you begin, have a USB boot device ready with the required Junos OS software install package as described in "Creating a USB Boot Device for a QFX5100 Switch" on page 174. The upgrade process overwrites the contents of the internal flash storage on the QFX5100 switch, so if desired, before proceeding with the upgrade, save existing configuration, backup configurations, and other stored files to a remote system, server, or other storage device. Then:

1. Insert the USB boot device into a USB port on the QFX5100 switch you need to upgrade, and power-cycle the QFX5100 switch.
2. The switch comes up, booting from the USB device and running the `Juniper Linux Installer`.

   > **NOTE**: If the switch doesn’t automatically boot from the USB device, press the ESC key while the switch reboots to bring up the BIOS boot manager so you can manually select to boot from the USB device.

   The installer menu prompts you to with the following options:

   ```
   Juniper Linux Installer - (c) Juniper Networks 2014
   Reboot
   Install Juniper Linux Platform
   Boot to host shell [debug]
   ```

   Tab through the options to select `Install Juniper Linux Platform`, and press Enter.
3. The installer displays status messages during the install process, creates and formats the local storage partitions, and installs the host OS and Junos OS software on the switch.

   After completing installation, the installer displays a list of boot options and after a few seconds automatically selects the default option to boot `Juniper Linux`. The switch boots the host OS and automatically selects and brings up Junos OS. Upon completion of the install and reboot process, the switch displays the Junos OS login prompt.

4. Log in to Junos OS on the switch, enter operational mode, and verify that the new version of software has been properly installed by running the show version command.

```
user@switch> show version
```

After the QFX5100 switch is running the “QFX 5e Series” software image, you can continue with any other configuration and setup to add the QFX5100 to a QFX5110 Virtual Chassis or VCF, such as configuring the Virtual Chassis ports (VCPs).

Release History Table

Release	Description
17.4R1	Starting in Junos OS Releases 17.3R2 and 17.4R1, you are not required to use the USB install method to upgrade a QFX5100 switch from a “QFX 5 Series” to a “QFX 5e Series” image. Instead, you can install the “QFX 5e Series” package directly using the CLI command *request system software add*

RELATED DOCUMENTATION

Understanding Mixed EX Series and QFX Series Virtual Chassis

Understanding Mixed Virtual Chassis Fabric

Configuring an EX4650 or a QFX Series Virtual Chassis

Understanding Virtual Chassis Fabric Configuration

Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis

IN THIS SECTION

Nonstop software upgrade (NSSU) enables you to upgrade the software running on all member switches in a *Virtual Chassis* with minimal network traffic disruption during the upgrade. This topic describes NSSU on EX Series and QFX Series Virtual Chassis that support this feature.

See these other references for information on using NSSU on the following specific platforms:

- EX8200 Virtual Chassis—For information on using NSSU with an EX8200 Virtual Chassis, see *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*.
- Virtual Chassis Fabric (VCF)—For information on using NSSU with a VCF, see *Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric*.

NOTE: Because NSSU upgrades the software on each Virtual Chassis member one at a time, upgrading using NSSU can take longer than an upgrade using the `request system software add` command.

You can reduce the amount of time an upgrade takes by configuring line-card upgrade groups on larger Virtual Chassis that support this feature. The Virtual Chassis upgrades the member switches in an upgrade group simultaneously, reducing the amount of time it takes to complete an upgrade. See *Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade*.

Benefits of NSSU

- No disruption to the control plane—NSSU uses *graceful Routing Engine switchover* (GRES) (and *nonstop active routing* (NSR) on applicable platforms) to ensure no disruption occurs to the control plane. During the upgrade process, the Virtual Chassis preserves interface, kernel, and routing protocol information.
- Minimal disruption to network traffic—NSSU minimizes network traffic disruption by upgrading member switches one at a time, enabling the primary and backup members to maintain their primary and backup roles (although primary role will change) without disrupting traffic, and permitting traffic to continue to flow through members in the linecard role that are not being upgraded.

Requirements for Performing an NSSU

Requirements for performing NSSU for a Virtual Chassis include:

- All Virtual Chassis members and all Routing Engines must be running the same Junos OS release.
- You must enable Graceful Routing Engine switchover (GRES).
- You must enable nonstop active routing (NSR) for applicable platforms.

  Although nonstop bridging (NSB) is not required to perform an NSSU, we also recommend enabling NSB before performing an NSSU on applicable platforms. NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly when the Routing Engine switches over during NSSU. See Configuring Nonstop Bridging on Switches (CLI Procedure).

- To minimize traffic disruption, you must configure link aggregation groups (LAGs) such that the member links of each LAG reside on different Virtual Chassis members, and configure Link Aggregation Control Protocol (LACP) to monitor LAG member link states. When one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG. For more information on configuring LAGs and LACP, see Configuring Link Aggregation and Configuring Aggregated Ethernet LACP (CLI Procedure).

  NOTE: When you upgrade an EX Series switch in a mixed Virtual Chassis to Junos OS Release 15.1 or later from a release earlier than Release 15.1, there might be a drop in traffic for up to 60 seconds.

  NOTE: During an NSSU operation, if you try to view LAG interface status on the primary Routing Engine member using the `show interfaces ae-`*`ae-interface-number`* CLI command, you might see incorrect or zero traffic counts. To work around this problem, run the command on the backup Routing Engine member instead if that member is already loaded and running.

Requirements for the Virtual Chassis or mixed Virtual Chassis members being upgraded using NSSU:

- The member switches must be connected in a ring topology so that no member is isolated as a result of another member being rebooted. This topology prevents the Virtual Chassis from splitting during an NSSU.
- The primary and backup member switches must be adjacent to each other in the ring topology. Adjacent placement ensures the primary and backup are always in sync while the member switches in linecard roles are rebooting.

- The Virtual Chassis is preprovisioned and you have explicitly assigned the linecard role to the member switches acting in a linecard role. The Virtual Chassis primary and backup member switches change primary role while one or the other is being upgraded during NSSU, but they must maintain their primary and backup routing engine roles, and the remaining switches must maintain their linecard roles.

- A two-member Virtual Chassis must have `no-split-detection` configured so that the Virtual Chassis doesn't split when an NSSU upgrades a member. See "Understanding Split and Merge in a Virtual Chassis" on page 71.

> **NOTE**: In an EX4300 Virtual Chassis running a Junos OS 13.2X50 release, you should enable the *vcp-no-hold-time* statement at the `[edit virtual-chassis]` hierarchy level before performing a software upgrade using NSSU. Without this option configured, the Virtual Chassis might split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you might have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see "Understanding Split and Merge in a Virtual Chassis" on page 71. This statement only affects EX4300 Virtual Chassis or mixed Virtual Chassis that contain EX4300 switches.

How an NSSU Works on a Virtual Chassis and Mixed Virtual Chassis

When you request an NSSU on a Virtual Chassis or mixed Virtual Chassis:

1. The Virtual Chassis primary verifies that:

   - The backup is online and running the same software version.
   - You enabled Graceful Routing Engine switchover (GRES), and, if applicable, nonstop active routing (NSR).
   - You used a preprovisioned configuration to set up the Virtual Chassis.

2. The primary copies the new software image to the backup and remaining linecard role members in sequence using `rcp`.

   (For QFX5100 Virtual Chassis only) Starting with Junos OS Release 14.1X53-D40, to optimize the time needed to complete an NSSU operation for a Virtual Chassis, the primary uses parallel `rcp` sessions to copy the new software to multiple members at a time (rather than waiting for the copy operation to complete to each member before starting to copy the software image to the next member). The primary uses a default algorithm to determine the number of parallel copy operations based on the number of members in the Virtual Chassis, or you can configure a specific number using the `rcp-count` configuration statement. See *rcp-count* for details.

> **NOTE**: If copying the new software to any member fails, NSSU terminates the upgrade process for the entire Virtual Chassis without rebooting any members, and logs the error condition. Starting with Junos OS Release 14.1X53-D40, if an NSSU copy operation to a member fails, the primary performs an additional error recovery measure to remove the new software from the members to which it was already transferred.

3. The primary restarts the backup member switch with the new software, and the backup resynchronizes with the primary.

4. The primary loads and reboots member switches that are in the linecard role, one at a time. The primary waits for each member to become online and active running the new software before rebooting the next member.

    - If you configured upgrade groups, the Virtual Chassis members in the first upgrade group load the new image and restart. When the members in that upgrade group are online again, the members in the next upgrade group load the new image and restart. (NSSU upgrades the groups in the order that they appear in the configuration.)

    - Traffic continues to flow through the other members during this process.

5. Rebooting continues until all active members have restarted with the new software.

> **NOTE**: If any linecard role member fails to reboot successfully, NSSU terminates the upgrade process and logs the error condition. In this case, to avoid Virtual Chassis instability, you should either back out the partial upgrade by restoring the old software and rebooting the members that were already rebooted with the new software, or try to manually reboot all members with the new software that was copied to them, so all members come online again running the same version of the software.
>
> Starting with Junos OS Release 14.1X53-D40, NSSU automatically invokes recovery measures if the reboot fails on any linecard role member, stopping the sequential reboot process and bringing down and rebooting the entire Virtual Chassis. The Virtual Chassis then cleanly brings up all members at the same time running the new software, which recovers Virtual Chassis stability more quickly than having an unstable Virtual Chassis running different versions of the software trying to converge.

6. After the primary has upgraded all members in the linecard role, it performs a graceful Routing Engine switchover and the upgraded backup member switch becomes the new primary.

7. The new primary upgrades the software on the original primary and automatically reboots it. After the original primary has rejoined the Virtual Chassis, you can optionally revert primary role to that switch by explicitly requesting another graceful Routing Engine switchover.

NSSU Limitations

You can't use NSSU to downgrade the software—that is, to install an earlier version of the software than is currently running on the switch. To install an earlier software version, use the `request system software add` command.

You can't roll back to the previous software version after you perform an upgrade using NSSU. If you need to roll back to the previous software version, you can reboot from the alternate root partition if you have not already copied the new software version into the alternate root partition.

NSSU and Junos OS Release Support

NSSU works only on some Virtual Chassis with particular *from* and *to* Junos OS Releases. Contact Juniper Networks Technical Assistance Center (JTAC) to confirm supported *from* and *to* releases if you are considering upgrading your Virtual Chassis using NSSU.

If your Virtual Chassis is running a software version that does not support NSSU or does not support the combination of *from* and *to* releases with NSSU, use the `request system software add` command to upgrade the member switches in the Virtual Chassis individually.

You can also refer to this network configuration example on how to manually upgrade a two-member QFX Series Virtual Chassis with minimal impact to traffic flow when NSSU is not supported:

- Two-Member QFX Series Virtual Chassis Upgrade Procedure

Overview of NSSU Configuration and Operation

For NSSU to succeed, the Virtual Chassis and member switches must meets the requirements in "Requirements for Performing an NSSU" on page 180. NSSU requires only those configuration steps.

If your Virtual Chassis meets the NSSU requirements, simply enter the `request system software nonstop-upgrade` command to start NSSU. See "Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade" on page 189 for details.

Release History Table

Release	Description
14.1X53-D40	(For QFX5100 Virtual Chassis only) Starting with Junos OS Release 14.1X53-D40, to optimize the time needed to complete an NSSU operation for a Virtual Chassis, the primary uses parallel `rcp` sessions to copy the new software to multiple members at a time (rather than waiting for the copy operation to complete to each member before starting to copy the software image to the next member).
14.1X53-D40	Starting with Junos OS Release 14.1X53-D40, if an NSSU copy operation to a member fails, the primary performs an additional error recovery measure to remove the new software from the members to which it was already transferred.
14.1X53-D40	Starting with Junos OS Release 14.1X53-D40, NSSU automatically invokes recovery measures if the reboot fails on any linecard role member, stopping the sequential reboot process and bringing down and rebooting the entire Virtual Chassis.

RELATED DOCUMENTATION

Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade | 189

Configuring Graceful Routing Engine Switchover in a Virtual Chassis

Understanding Aggregated Ethernet Interfaces and LACP for Switches

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

IN THIS SECTION

You can configure line-card upgrade groups for nonstop software upgrade (NSSU) operations on supporting platforms. Line-card upgrade groups can reduce the total time required to complete an NSSU operation and enable you to control the upgrade sequence among the switches being upgraded.

How Line-card Upgrade Groups Work with Nonstop Software Upgrade

With NSSU, you can upgrade software on supporting switches with redundant Routing Engines, a Virtual Chassis, or a Virtual Chassis Fabric (VCF) using a single command with minimal disruption to network traffic.

In its default configuration, NSSU upgrades each line card in a switch or linecard role member in a Virtual Chassis or VCF one at a time. Traffic continues to flow through the other line cards or members while each one is being restarted as part of the upgrade. This behavior minimizes traffic disruption if you configure link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards or members. As a result, when one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.

When you configure line-card upgrade groups for NSSU, NSSU upgrades all of the devices in each upgrade group at the same time instead of sequentially, reducing the total time needed to complete the upgrade on all line cards or members.

To achieve minimal traffic disruption during an NSSU operation, you must define the line-card upgrade groups such that the member links of the LAGs reside on line cards or members that are in different upgrade groups. For information on how to configure LAGs, see Configuring Aggregated Ethernet Links (CLI Procedure).

NSSU upgrades the groups in the order that they appear in the configuration (in other words, in the order you configure them). As a result, you can also define upgrade groups to control the upgrade sequence during an NSSU operation.

To configure upgrade groups, use the `upgrade-group` configuration statement in the `[edit chassis nssu]` hierarchy.

Line-card Upgrade Groups Support

The following platforms support NSSU line-card upgrade groups:

- EX4650 Virtual Chassis with more than three member switches
- QFX3500, QFX3600, and QFX5100 Virtual Chassis
- QFX5100 Virtual Chassis Fabric (VCF)
- EX6200 or EX8200 switches with redundant Routing Engines
- EX8200 Virtual Chassis

Configure Line-Card Upgrade Groups on an EX4650 Virtual Chassis, a QFX Series Virtual Chassis or a QFX5100 VCF

When you configure line-card upgrade groups on an EX4650 Virtual Chassis, a QFX Series Virtual Chassis, or a QFX5100 VCF, whose switches do not have separate line cards, you use only the `fpcs` option to specify the Virtual Chassis or VCF member IDs that you want to include in an upgrade group. You don't need to use the `member` option.

- To create an upgrade group and add a Virtual Chassis or VCF member switch to the upgrade group, configure the upgrade group name and specify the member number using the `fpcs` option:

  ```
  [edit chassis]
  user@switch# set nssu upgrade-group group-name fpcs member-number
  ```

  For example, to create an upgrade group called `vcf` and add linecard role member 2 to that group:

  ```
  [edit chassis]
  user@switch# set nssu upgrade-group vcf fpcs 2
  ```

  If `vcf` already exists, this command adds member 2 to `vcf`.

- To create an upgrade group that contains multiple members in a Virtual Chassis or VCF, specify multiple member numbers enclosed in square brackets after the `fpcs` option:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs [list-of-member-numbers]
```

  For example, to create an upgrade group called `vc1` that contains members 1 and 2:

```
[edit chassis]
user@switch# set nssu upgrade-group vc1 fpcs [1 2]
```

Make sure you commit the configuration before starting an NSSU operation.

Configure Line-Card Upgrade Groups on Standalone EX6200 or EX8200 Switches

To configure line-card upgrade groups on a standalone EX6200 or EX8200 switch:

- To create an upgrade group and add a line card to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs slot-number
```

  For example, to create an upgrade group called `group3` and add the line card in slot 5 to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group3 fpcs 5
```

  If `group3` already exists, this command adds line card 5 to `group3`.

- To create an upgrade group and add multiple line cards to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs [list-of-slot-numbers]
```

For example, to create an upgrade group called `primary` and add line cards in slots 1, 4, and 7 to it:

```
[edit chassis]
user@switch# set nssu upgrade-group primary fpcs [1 4 7]
```

If `primary` already exists, this command adds line cards in slots 1, 4, and 7 to `primary`.

SEE ALSO

Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)

Configure Line-Card Upgrade Groups on an EX8200 Virtual Chassis

To configure line-card upgrade groups on an EX8200 Virtual Chassis:

- To create an upgrade group and add a line card on a Virtual Chassis member to it:

  ```
  [edit chassis]
  user@switch# set nssu upgrade-group group-name member member-id fpcs slot-number
  ```

  For example, to create an upgrade group called `primary-ny` and add the line card on member 1 in slot 5 to it:

  ```
  [edit chassis]
  user@switch# set nssu upgrade-group primary-ny member 1 fpcs 5
  ```

  If `primary-ny` already exists, this command adds line card 5 on member 1 to `primary-ny`.

- To create an upgrade group that contains multiple line cards on a Virtual Chassis member:

  ```
  [edit chassis]
  user@switch# set nssu upgrade-group group-name member member-id fpcs [list-of-slot-numbers]
  ```

For example, to create an upgrade group called `primary-ny` that contains the line cards in slots 1 and 2 on member 0 and in slots 3 and 4 on member 1:

```
[edit chassis]
user@switch# set nssu upgrade-group primary-ny member 0 fpcs [1 2]

[edit chassis]
user@switch# set nssu upgrade-group primary-ny member 1 fpcs [3 4]
```

SEE ALSO

Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

RELATED DOCUMENTATION

Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis

Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade

Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric

Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade

Understanding Nonstop Software Upgrade on EX Series Switches

Upgrading Software Using Nonstop Software Upgrade on EX Series Virtual Chassis and Mixed Virtual Chassis (CLI Procedure)

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade

IN THIS SECTION

Nonstop software upgrade (NSSU) enables you to upgrade the software running on all member switches of supported Virtual Chassis with minimal traffic disruption during the upgrade.

> **NOTE**: NSSU works only on some Virtual Chassis with certain *from* and *to* Junos OS Releases. Use the `request system software add` command to upgrade the member switches in the Virtual Chassis individually if the Virtual Chassis is running a software version that does not support NSSU or does not support the combination of *from* and *to* releases.
>
> You can also refer to Two-Member QFX Series Virtual Chassis Upgrade Procedure, a network configuration example on how to manually upgrade a two-member QFX Series Virtual Chassis with minimal impact to traffic flow when NSSU is not supported.

Preparing the Switch for Software Installation

Before you begin installing the new software using NSSU:

- Ensure that the Virtual Chassis is connected and configured correctly to support the NSSU process. See "Requirements for Performing an NSSU" on page 180.
- Verify that the members are running the same version of the software:

  ```
  user@switch> show version
  ```

  If the Virtual Chassis or mixed Virtual Chassis members are not running the same version of the software, use the `request system software add` command to upgrade the software on the inconsistent members.

- Ensure that graceful Routing Engine switchover (GRES) is enabled, or for applicable platforms, make sure nonstop active routing (NSR) is enabled, which also enables graceful Routing Engine switchover. See Configuring Nonstop Active Routing on Switches for more information.

  To check the nonstop active routing state to verify both NSR and GRES are enabled:

  ```
  user@switch> show task replication
  ```

- (Optional for applicable platforms) Enable nonstop bridging (NSB), which ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU. See Configuring Nonstop Bridging on Switches (CLI Procedure) for details.

- For a two-member Virtual Chassis, make sure you configured `no-split-detection` so the Virtual Chassis does not split when NSSU upgrades one of the members. See "Disabling Split and Merge in a Virtual Chassis" on page 140.

- In an EX4300 Virtual Chassis running a Junos OS 13.2X50 release, you should set the `vcp-no-hold-time` option at the `[edit virtual-chassis]` hierarchy level before performing a software upgrade using NSSU, otherwise the Virtual Chassis might split during the upgrade. A split Virtual Chassis can disrupt your network, and you might need to manually reconfigure your Virtual Chassis after the NSSU if the split-and-merge feature was disabled. For more information about a Virtual Chassis split, see "Understanding Split and Merge in a Virtual Chassis" on page 71. This statement only affects EX4300 Virtual Chassis or mixed Virtual Chassis that contain EX4300 switches.

  To configure this option:

  ```
  user@switch# set virtual-chassis vcp-no-hold-time
  ```

- On a QFX5100 Virtual Chassis with line-card upgrade groups configured, you should enable the `lc-reboot-delay` option to configure a delay for when adjacent members in a line card group reboot. Without this option, when the next member reboots, approximately two minutes after the previous member reboots and joins the Virtual Chassis, the previous rebooted member might not be ready to carry traffic. This delay helps prevent dropping traffic when there are two adjacent line card members with interfaces that are part of a common link aggregation group (LAG).

  We recommend setting a 200-second delay (the allowable range is 0 to 600 seconds). To configure this delay:

  ```
  [edit chassis]
  user@switch# set chassis nssu lc-reboot-delay 200
  ```

- (Optional) Back up the system software (Junos OS, the active configuration, and log files) on each member to an external storage device as desired using the `request system snapshot` command.

Upgrading the Software Using NSSU

This procedure describes how to upgrade the software running on all Virtual Chassis or mixed Virtual Chassis members using NSSU. When the upgrade completes, all members are running the new version of the software. The upgrade includes a graceful Routing Engine switchover, so the original Virtual Chassis backup member switch becomes the new primary.

During NSSU, the primary copies the new software image to all the members in the Virtual Chassis and reboots them in turn. If copying the new software to a member fails or rebooting a member fails, NSSU terminates the upgrade process and logs the error. In this case, you must manually perform recovery measures for members left in an incompatible state to restore all members to running the same version of the software. Starting in Junos OS Release 14.1X53-D40, NSSU automatically invokes recovery measures after either of these failures, as follows:

- if NSSU terminates due to a copy error, the primary removes the new image from any members to which it was already copied.
- If any member fails to reboot, NSSU automatically initiates a clean Virtual Chassis restart by bringing down and rebooting the entire Virtual Chassis. All members come up running the new software at the same time. This action cleanly recovers correct Virtual Chassis operation more quickly than having an unstable Virtual Chassis running different versions of the software trying to converge.

> **NOTE**: Junos OS software images with enhanced automation are only supported on a non-mixed Virtual Chassis with QFX5100 switches. Also, you can't perform an NSSU from a standard Junos OS software image to a Junos OS software image with enhanced automation, or from a Junos OS software image with enhanced automation to a standard Junos OS software image.

To upgrade all members in a Virtual Chassis using NSSU:

1. Download the software package as described in Installing Software Packages on QFX Series Devices. If you are upgrading a mixed Virtual Chassis, download the software packages for the different switch types.
2. Copy the software package or packages to the Virtual Chassis. We recommend that you copy the file or files to the `/var/tmp` directory on the primary.
3. Use the console connection or the virtual management Ethernet (VME) interface to log in to the Virtual Chassis or mixed Virtual Chassis. You can monitor the progress of the primary switch reboot if you use a console connection.
4. Start the NSSU:
   - On a Virtual Chassis where all members use the same software image, enter:

     ```
     user@switch> request system software nonstop-upgrade force-host /var/tmp/package-name.tgz
     ```

     where *package-name*`.tgz` is the software package name, for example, `jinstall-qfx-3-13.2X50-D15.3-domestic-signed.tgz`.

- On a mixed Virtual Chassis where members might use different software images, enter the `request system software nonstop-upgrade` command with the `set` option to specify more than one software package name:

  ```
  user@switch> request system software nonstop-upgrade set [/var/tmp/package-
  name1.tgz /var/tmp/package-name2.tgz]
  ```

  For example, `/var/tmp/package-name1.tgz` and `/var/tmp/package-name2.tgz` might specify software packages for EX4200 and EX4500 switches in a mixed EX Series Virtual Chassis with EX4200 and EX4500 switches.

The switch displays status messages similar to the following messages as the upgrade executes:

```


Chassis ISSU Check Done
NSSU: Validating Image
NSSU: Preparing Backup RE
Installing image on other FPC's along with the backup

Checking pending install on fpc1
Pushing bundle to fpc1
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc1

Checking pending install on fpc2
Pushing bundle to fpc2
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc2

Rebooting fpc1
NSSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis NSSU Started
NSSU: Preparing Daemons
NSSU: Daemons Ready for NSSU
NSSU: Starting Upgrade for FRUs
NSSU: Preparing for Switchover
NSSU: Ready for Switchover
```

```
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
relinquish mastership
NSSU: IDLE

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY



Shutdown NOW!
[pid 9336]
```

5. Log in after the original primary switch reboot completes. To verify that the software is upgraded on all Routing Engines in the Virtual Chassis, enter the following command:

```
user@switch> show version
```

6. To ensure the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partitions of all members:

```
user@switch> request system snapshot slice alternate all-members
```

With resilient dual-root partitions, the switch can boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

NOTE: After an upgrade is complete, please verify `syslog`, `show chassis fabric errors`, `show chassis fabric fpcs`, and `show system alarms`.

If the FPCs or fabric display any errors, set alarms for specific errors. Configure `pfe-offline` as error action to mitigate outages.

RELATED DOCUMENTATION

Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis | **178**

Configuring Dual-Root Partitions

5

CHAPTER

Configuration Statements

aliases (Virtual Chassis)

IN THIS SECTION

Syntax

```
aliases {
    serial-number serial-number {
        alias-name alias-name;
    }
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Optionally create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). You can more clearly identify the member switches in your Virtual Chassis or VCF by assigning text labels (alias names) to them, associating the alias name with member switch serial numbers. Aliases are for administrative purposes only and have no effect on member switch operations.

The `show virtual-chassis` command displays assigned alias names in the `Alias-Name` output field.

Options

serial-number The permanent serial number for the member switch in the Virtual Chassis or VCF that you want to label with an alias name. Serial number values are case-sensitive.

The `show virtual-chassis` command displays the serial number for all members in the Virtual Chassis or VCF in the `Serial No` output field.

alias-name The text label you want to assign to the member switch.

For example, if you configure alias name `dc-floor-1` in a VCF for member switch 0 with serial number AB012345678 as follows:

```
[edit virtual-chassis aliases]
user@switch# set serial-number AB0123456789 alias-name dc-floor-1
```

then you'll see the alias **dc-floor-1** in the Alias-Name column in the output from the `show virtual-chassis` command:

show virtual-chassis

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 9d5d.5556.919a
Fabric Mode: Enabled
                                                                  Mstr
Member ID  Status   Serial No      Alias-Name     Model           prio    Role
0 (FPC 0)  Prsnt    AB0123456789   dc-floor-1     qfx5100-48s-6q  129     Master
.
.
.
```

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Configuring an EX4650 or a QFX Series Virtual Chassis

Understanding Virtual Chassis Fabric Components

Understanding Virtual Chassis Components

auto-sw-update

IN THIS SECTION

Syntax

```
auto-sw-update {
    (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
    package-name package-name;
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable the automatic software update feature for Virtual Chassis or Virtual Chassis Fabric (VCF) configurations.

You should only use the keywords that specify a device or device family—`ex-4300`, `ex-4600`, `qfx-3`, and `qfx-5` —when configuring automatic software update on a mixed mode Virtual Chassis or Virtual Chassis Fabric (VCF). You can simply specify the *package-name* without specifying individual device keywords in non-mixed Virtual Chassis or VCF topologies.

NOTE: The automatic software update feature is not supported for a mixed mode EX4300 Virtual Chassis with a combination of EX4300 multigigabit model (EX4300-48MP) switches and other EX4300 switches.

You must enter the `auto-sw-update` statement multiple times—once for each device family in your mixed Virtual Chassis or VCF—in most scenarios when enabling the automatic software update for a mixed Virtual Chassis or VCF.

The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. You do not, therefore, need to specify the `ex-4500` keyword when configuring automatic software update for a mixed Virtual Chassis that include EX4500 and EX4550 switches only. You also only have to enter the `ex-4500` keyword once to configure automatic software update for all EX4500 and EX4550 member switches in the s37628)ame mixed Virtual Chassis.

The Junos OS package for a QFX3500 device updates the software for QFX3500 and QFX3600 devices. You do not, therefore, need to specify the `qfx-3` keyword when configuring automatic software update for a Virtual Chassis composed entirely of QFX3500 and QFX3600 devices. You also have to enter the `qfx-3` keyword only once to configure automatic software update for all QFX3500 and QFX3600 member devices in the same mixed Virtual Chassis.

The Junos OS package for a QFX5110 device updates the software for both QFX5110 and QFX5100 devices in a QFX5110 Virtual Chassis or VCF, so you do not need to specify the `qfx-5` keyword when configuring automatic software update for a QFX5110 Virtual Chassis or VCF composed of both types of devices.

> **CAUTION**: A QFX5100 switch running a Junos OS software image with “-qfx-5-” in the package filename *must* first be upgraded to a Junos OS software image with “-qfx-5e-” in the package filename before it can be added to a mixed QFX5110 Virtual Chassis or VCF. The automatic software update process cannot update a switch from a “-qfx-5-” image to a “-qfx-5e-” image. See *Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric*.
>
> After a QFX5100 switch is installed with a “-qfx-5e-” Junos OS software image, the automatic software update process can successfully update the switch automatically with a different version of a “-qfx-5e-” Junos OS image to match the other members in the Virtual Chassis or VCF.

The remaining statement is explained separately. See CLI Explorer.

Default

The automatic software update feature is disabled.

Options

package-name ***package-name*** — Specify the software package name or a path to a Junos OS software image to update new members being added to a Virtual Chassis or VCF. .

- If the software package is located on a local directory on the switch, use the following format for *package-name*:

  /pathname/package-name

- If the software package is to be downloaded and installed from a remote location, use one of the following file access and transfer formats:

  ftp://*hostname*/*pathname*/*package-name*

  ftp://*username*:*prompt*@ftp.*hostname.net*/*package-name*

  http://*hostname*/*pathname*/*package-name*

ex-4200 Specify a path to a Junos OS image for an EX4200 switch when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

ex-4300 Specify a path to a Junos OS image for an EX4300 switch when enabling automatic software update for a mixed Virtual Chassis or VCF. Automatic software update is not supported for mixed mode EX4300 Virtual Chassis with a combination of EX4300 multigigabit model (EX4300-48MP) switches and other EX4300 switches, and this option is not applicable in that case.

ex-4500 Specify a path to a Junos OS image for an EX4500 switch, an EX4550 switch, or both types of switches when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. Therefore, you only enter this command once to upgrade the EX4500 and EX4550 member switches in the same mixed Virtual Chassis.

The `ex-4500` keyword also does not need to be specified when configuring automatic software update for a mixed EX4500 and EX4550 Virtual Chassis.

ex-4600 Specify a path to a Junos OS image for an EX4600 switch when enabling automatic software update for a mixed Virtual Chassis.

qfx-3 Specify a path to a Junos OS image for a QFX3500, QFX3600, or both types of devices when enabling automatic software update for a mixed Virtual Chassis or VCF.

qfx-5 Specify a path to a Junos OS image for a QFX5100 device when enabling automatic software update for a mixed QFX5100 Virtual Chassis or VCF. You do not need to use this option for a QFX5110 Virtual Chassis or VCF that contains both QFX5110 and QFX5100 switches because the same software image runs on either type of switch.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statements introduced in Junos OS Release 10.0 for EX Series switches.

The `ex-4200` and `ex-4500` options introduced in Junos OS Release 12.2 for EX Series switches.

Statements introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

The `ex-4300`, `qfx-3`, and `qfx-5` options introduced in Junos OS Release 13.2X51-D20.

Statements introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

The `ex-4600` option introduced in Junos OS Release 13.2X51-D25.

RELATED DOCUMENTATION

Configuring Automatic Software Update on Virtual Chassis Member Switches

Understanding Software Upgrades in a Virtual Chassis

Understanding Software Upgrades in a Virtual Chassis Fabric

Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches

fpcs (NSSU Upgrade Groups)

IN THIS SECTION

Syntax

```
fpcs (slot-number | [list-of-slot-numbers]);
```

Hierarchy Level

```
[edit chassis nssu upgrade-group group-name],
[edit chassis nssu upgrade-group group-name member member-id]
```

Description

Configure switch line cards, Virtual Chassis member switches, or Virtual Chassis Fabric (VCF) member switches as part of an NSSU upgrade group.

To reduce the time an NSSU takes, you can configure line-card upgrade groups for an EX6200 or EX8200 switch with redundant Routing Engines; an EX8200 Virtual Chassis; an EX4650 Virtual Chassis with more than three member switches; QFX3500, QFX3600, and QFX5100 Virtual Chassis; or a QFX5100 Virtual Chassis Fabric (VCF). NSSU upgrades the devices in the order in which you configure the upgrade groups, so you can also use upgrade groups to control the upgrade sequence.

For switches that have separate line cards, use this statement to assign one or more line cards to an NSSU upgrade group based on their line-card slot numbers.

For Virtual Chassis or VCF member switches that do not have separate line cards, use this statement to assign one or more Virtual Chassis or VCF members to an NSSU upgrade group by specifying their member IDs.

NOTE: For a Virtual Chassis or VCF, you do not use this statement with the `member` option. When to use the `member` statement hierarchy is explained next.

To configure an upgrade group that includes line cards on switches that support multiple line cards and comprise a Virtual Chassis, use this statement with the `member` option to specify the Virtual Chassis member ID and the desired line card slot number or numbers on that member switch to include in the upgrade group. Use multiple statements to add line cards from different Virtual Chassis members to the upgrade group.

Options

list-of-slot-numbers A list of slot numbers of multiple line cards or member IDs of Virtual Chassis or VCF members to be included in the upgrade group. Separate multiple slot numbers or member IDs with spaces and enclose the list in square brackets—for example: [3 4 7].

slot-number The slot number of a single line card or member ID of a Virtual Chassis or VCF member to be included in the upgrade group.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

id

IN THIS SECTION

Syntax

```
id id;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure the alphanumeric string that identifies a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.

Options

id—Virtual Chassis ID (VCID), which uses the ISO family address format—for example, **9622.6ac8.5345**.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

- Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge
- Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge
- Configuring an EX4650 or a QFX Series Virtual Chassis
- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring an EX8200 Virtual Chassis (CLI Procedure)
- Understanding Virtual Chassis Member ID Numbering in an EX8200 Virtual Chassis

location (Virtual Chassis)

IN THIS SECTION

Syntax

```
location location;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Set a description of the location of the Virtual Chassis or VCF member switch or external Routing Engine.

The `Location` field is visible to users who enter the `show virtual-chassis status detail` command.

Setting this description has no effect on the operation of the member device.

Options

location—Location of the current member switch or external Routing Engine. The `location` can be any single word.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Configuring an EX4650 or a QFX Series Virtual Chassis

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

mac-persistence-timer

IN THIS SECTION

Syntax

```
mac-persistence-timer (minutes | disable);
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Specify how long the Virtual Chassis or VCF continues to use the MAC address of the switch that was originally configured in the primary role as the system MAC base address after the original primary switch is removed from the Virtual Chassis or VCF. The system MAC base address does not change in the event of a switchover provided the switch originally configured in the primary role remains a member of the Virtual Chassis or VCF.

The maximum timer limit is 60 minutes starting in Junos OS Release 12.2. There are no minimum or maximum timer limits in prior Junos OS releases.

Default

The MAC persistence timer is set to 10 minutes by default.

Options

minutes Time in minutes that the member switch in the backup role continues to use the system MAC base address of the old primary before using its own system MAC base address after the switch in the primary role is physically disconnected or removed from the Virtual Chassis or VCF.

`disable` Disable the MAC persistence timer. The system MAC base address never changes when the MAC persistence timer is disabled, even when the switch in the primary role is physically disconnected or removed from the Virtual Chassis or VCF.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Option `disable` introduced in Junos OS Release 12.2 for EX Series switches.

The maximum timer limit changed from no maximum timer limit to 60 minutes in Junos OS Release 12.2 for EX Series switches.

RELATED DOCUMENTATION

Configuring the Timer for the Backup Member to Start Using Its Own MAC Address as Primary of a Virtual Chassis

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

mastership-priority

IN THIS SECTION

Syntax

```
mastership-priority number;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Assign a primary-role priority value to a member of a Virtual Chassis or VCF. The primary-role priority is the most important factor that determines which member assumes the primary Routing Engine role in the Virtual Chassis or VCF.

A Virtual Chassis or VCF configuration has two member switches running in the Routing Engine role, one as the primary and one as the backup. The remaining members operate in the linecard role. In a preprovisioned Virtual Chassis or VCF, you configure the *role* for each member by associating it with the member switch's serial number, and the Virtual Chassis assigns default primary-role priorities based on the assigned role. In a non-provisioned Virtual Chassis or VCF, you can configure roles explicitly, or the primary-role priorities of the members implicitly designate the role of each member based on the primary-role election algorithm.

The Virtual Chassis or VCF elects the member switch with the highest primary-role priority value as the primary Routing Engine. When two members have the same highest primary-role priority value among all the members, the primary-role election algorithm uses other factors to determine which of the two members functions as the primary and which is the backup. See Understanding How the Primary in a Virtual Chassis Is Elected for all the conditions considered in primary-role election.

You should assign the *same* (highest) primary-role priority value to the members you want to be the primary and backup Routing Engine members to ensure reliable graceful Routing Engine switchover (GRES) operation. This and other primary-role election considerations also help keep primary role from switching back and forth rapidly between the two members under failover conditions.

This statement is not used for the EX8200 Virtual Chassis, which determines primary role by external Routing Engine uptime.

A switch with a primary-role priority of `0` never takes the primary or backup role.

Default

128

Options

number—Primary-role priority value.

- **Range:** 0 through 255

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric
Configuring an EX4650 or a QFX Series Virtual Chassis
Configuring an EX2300, EX3400, or EX4300 Virtual Chassis
Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

member

IN THIS SECTION

Syntax

```
member member-id {
    fabric-tree-root;
    location location;
    mastership-priority number;
    no-management-vlan;
    serial-number serial-number;
    role role;
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Configure a switch or an XRE200 External Routing Engine as a member of a Virtual Chassis or a Virtual Chassis Fabric (VCF) with characteristics specified by the available options.

Default

When an EX Series switch or a QFX Series devices configured in standalone mode is powered on but not interconnected through its Virtual Chassis ports (VCPs) with other member switches, its default member ID is `0`.

There is no default member ID in an EX8200 or EX9200 Virtual Chassis. An EX8200 or EX9200 Virtual Chassis must be preprovisioned, and that process configures the member IDs.

Options

member-id—Identifies a specific member switch of a Virtual Chassis or VCF configuration.

The exact range for a specific Virtual Chassis or VCF depends on the number of switches allowed in the Virtual Chassis or VCF.

In an EX8200 Virtual Chassis, member IDs 0 through 7 are reserved for EX8200 member switches and member IDs 8 and 9 are reserved for the primary and backup external Routing Engines.

The remaining statement options set characteristics of the Virtual Chassis or VCF member, and are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric
Preprovisioning a Virtual Chassis Fabric
Adding a Device to a Virtual Chassis Fabric
Configuring an EX4650 or a QFX Series Virtual Chassis
Configuring an EX2300, EX3400, or EX4300 Virtual Chassis
Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis
Configuring an EX9200 Virtual Chassis
Configuring an EX4650 or a QFX Series Virtual Chassis

member (NSSU Upgrade Groups)

IN THIS SECTION

Syntax

```
member member-id {
    fpcs (slot-number | [list-of-slot-numbers]);
}
```

Hierarchy Level

```
[edit chassis nssu upgrade-group group-name]
```

Description

Specify the Virtual Chassis member whose line-card slot numbers you are assigning to an NSSU upgrade group.

> **NOTE**: This statement is not applicable to Virtual Chassis or VCF member switches that do not support separate line cards. To configure Virtual Chassis or VCF member switches that do not have separate line cards into an NSSU upgrade group, use the `fpcs` statement alone, and specify the Virtual Chassis or VCF member IDs to include in the upgrade group in place of line card slot numbers.

To reduce the time an NSSU takes, you can configure NSSU line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines; EX8200 Virtual Chassis; QFX3500, QFX3600, and QFX5100 Virtual Chassis; and Virtual Chassis Fabric (VCF).

To configure an upgrade group that includes line cards on different switches that support multiple line cards and comprise a Virtual Chassis, use this statement hierarchy with the `fpcs` option to first specify the Virtual Chassis member ID and then desired line card slot number or numbers on that member switch to include in the upgrade group. Use multiple statements to add line cards from different Virtual Chassis members to the upgrade group.

Options

member-id The ID of the Virtual Chassis or VCF member switch containing one or more line cards to include in an NSSU upgrade group.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

no-auto-conversion

IN THIS SECTION

Syntax

```
no-auto-conversion;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Disable automatic Virtual Chassis port (VCP) conversion in a Virtual Chassis.

> **NOTE**: This statement is available as of Junos OS Releases 14.1X53-D47, 17.4R2, 18.1R3, 18.2R2, and 18.3R1 only on EX Series and QFX Series switches in a Virtual Chassis that have automatic VCP conversion enabled by default, which include all EX4300, EX4600, EX4650, and QFX Series switches that support Virtual Chassis. See "Automatic Virtual Chassis Port (VCP) Conversion" on page 45 for details.

Configuring this statement disables the feature for the entire Virtual Chassis, and subsequently removing this statement returns the Virtual Chassis to the default behavior with automatic VCP conversion enabled.

VCP links connect members in a Virtual Chassis using supported VCP ports on the member switches. Ports that are supported as VCPs and are not dedicated VCPs or configured as VCPs by default must be converted into VCPs to use them to interconnect Virtual Chassis members. You can manually convert a supported port into a VCP port using the `request virtual-chassis vc-port` command, or when automatic VCP conversion is enabled, a port will be automatically converted into a VCP under the following conditions:

- LLDP is enabled on the interfaces for the members on both sides of the link. The two interfaces exchange LLDP packets to accomplish the port conversion.
- The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the `set virtual-chassis member` command.
- The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

If you want to deterministically control VCP port conversion during a preprovisioned Virtual Chassis configuration or expansion, you might want to disable this feature and set up all VCPs manually. Otherwise, you can use the default automatic VCP conversion behavior in a “plug and play” approach to simplify adding a new switch to an existing Virtual Chassis or adding a redundant VCP link between two existing members of a Virtual Chassis.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Releases 14.1X53-D47, 17.4R2, 18.1R3, 18.2R2, and 18.3R1.

RELATED DOCUMENTATION

Understanding Virtual Chassis Components | 35
Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port | 133
Adding a New Switch to an Existing EX2300, EX3400, EX4300, or EX4400 Virtual Chassis | 108
Adding an EX4600 Switch to a Mixed or Non-mixed Virtual Chassis | 113
Adding a New Switch to an Existing EX4650 or QFX Series Virtual Chassis | 116

no-management-vlan

IN THIS SECTION

Syntax

```
no-management-vlan;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Remove the specified member’s out-of-band management port from the virtual management Ethernet (VME) global management VLAN of the Virtual Chassis or VCF configuration.

For a member that is functioning in a linecard role, you can use this configuration to reserve the member's management Ethernet port for local troubleshooting:

```
virtual-chassis {
    member 2 {
```

```
        no-management-vlan;
        }
    }
```

You cannot configure the IP address for a local management Ethernet port using the CLI or the J-Web interface. To do this, you need to use the shell `ifconfig` command.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Understanding Global Management of a Virtual Chassis

Understanding Virtual Chassis Fabric Configuration

no-split-detection

IN THIS SECTION

Syntax

```
no-split-detection;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Disable the split and merge feature in a Virtual Chassis or VCF configuration.

The split and merge feature is enabled by default when you initially set up a Virtual Chassis. If you have only two members in your Virtual Chassis, we strongly recommend that you configure `no-split-detection` to disable the split and merge feature. This feature makes sure both switches remain in the correct Virtual Chassis roles in the event of a Virtual Chassis split. If you expand the two-member Virtual Chassis later to add more members, delete the `no-split-detection` configuration item to re-enable the split and merge feature again.

Default

The split and merge feature is enabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Understanding Split and Merge in a Virtual Chassis | 71

Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge

Disabling Split and Merge in a Virtual Chassis | 140

Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge | 144

nssu

IN THIS SECTION

Syntax

```
nssu {
    rcp-count number;
    upgrade-group group-name {
```

```
        fpcs (slot-number | [list-of-slot-numbers]);
        member member-id {
            fpcs (slot-number | [list-of-slot-numbers]);
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure parameters that affect the nonstop software upgrade (NSSU) process.

> **NOTE**: You use the *request system software nonstop-upgrade* command to initiate NSSU.

The `rcp-count` option (available only on QFX5100 switches) sets the number of parallel `rcp` sessions that NSSU uses to copy the new software to multiple Virtual Chassis or VCF member switches at a time.

The `upgrade-group` options define line-card upgrade groups for NSSU. When you initiate NSSU with at least one upgrade group configured, NSSU upgrades the line cards or Virtual Chassis or VCF members in each upgrade group to the new software version at the same time, in the order in which you configured them. Upgrade groups reduce the time required to complete an NSSU operation and control the order in which the line cards or members are upgraded.

Line-card upgrade groups are supported on some EX Series switches and EX Series Virtual Chassis that support NSSU and on a QFX5100 VCF.

These statements are all explained separately. You can also consult CLI Explorer.

Default

If you do not configure `rcp-count`, NSSU uses a default algorithm to determine the number of parallel `rcp` sessions to use based on the number of members in the Virtual Chassis or VCF.

If you do not define any line-card upgrade groups, NSSU upgrades line cards or members of a Virtual Chassis or VCF one at a time in ascending order by slot or member number.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`rcp-count` statement introduced in Junos OS Release 14.1X53-D40 for QFX5100 switches only.

RELATED DOCUMENTATION

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

preprovisioned

IN THIS SECTION

Syntax

```
preprovisioned;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable the preprovisioned configuration mode for a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.

When the preprovisioned configuration mode is enabled, you cannot use the CLI or the J-Web interface to change the primary-role priority or member ID of member switches.

You must use this statement to configure an EX8200 Virtual Chassis. Nonprovisioned configuration of an EX8200 Virtual Chassis is not supported.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Preprovisioning a Virtual Chassis Fabric

Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File

Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines

Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)

Configuring an EX8200 Virtual Chassis (CLI Procedure)

Configuring an EX9200 Virtual Chassis

Configuring an EX4650 or a QFX Series Virtual Chassis

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

rcp-count

IN THIS SECTION

Syntax

```
rcp-count number;
```

Hierarchy Level

```
[edit chassis nssu]
```

Description

(QFX5100 Virtual Chassis or Virtual Chassis Fabric [VCF] only) Optionally override the default algorithm that defines the number of parallel `rcp` sessions used for copying the new software image to member switches in a Virtual Chassis or VCF for nonstop software upgrade (NSSU).

You can upgrade the software running on all members of a Virtual Chassis or VCF using NSSU. At the beginning of the upgrade process, NSSU uses `rcp` to copy the new software from the primary switch to each of the member switches. Starting in Junos OS Release 14.1X53-D40, to minimize the time to copy the files to all members, NSSU uses parallel `rcp` sessions to copy the software to multiple members at the same time, rather than copying sequentially to each member in turn.

By default, the number of parallel `rcp` sessions that NSSU launches at a time follows an algorithm that optimizes the transfer time based on the total number of members in the Virtual Chassis or VCF, so that the new software is transferred to about half the members in parallel up to a maximum of 8 members at one time. After the first set of parallel copy operations are complete, NSSU launches the next set of parallel copy operations, and so on as needed until all members are updated.

Although the default value is usually optimal, you can configure the `rcp-count` statement to set a specific value for the number of parallel copy sessions instead. For example, you might want to use a lower number of parallel sessions to avoid potential impact on Virtual Chassis or VCF forwarding activity during NSSU.

After copying the new software to all members using as many parallel `rcp` sessions as needed, the NSSU process continues by rebooting each member with the new software in sequence, starting with the member in the backup Routing Engine role.

NOTE: If copying the new software to any member fails, NSSU terminates the upgrade process for the entire Virtual Chassis or VCF, logs the error condition, and initiates an error recovery measure to remove the new software from the members to which it was already transferred.

Options

number Number of parallel `rcp` sessions for NSSU to launch at a time, which overrides the default algorithm.

Allowable values are 0 through 8. Specifying a value of 0 suppresses parallel copy sessions; the new software is copied to each member sequentially.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1X53-D40.

RELATED DOCUMENTATION

Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis

Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric

role

IN THIS SECTION

Syntax

```
role (line-card | routing-engine);
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

Specify the roles of the members of the Virtual Chassis or a Virtual Chassis Fabric (VCF) in a preprovisioned Virtual Chassis. For a mixed Virtual Chassis or VCF, see Understanding Mixed EX Series and QFX Series Virtual Chassis or *Understanding Mixed Virtual Chassis Fabric* for any recommendations or requirements for assigning the Routing Engine role based on the types of switches comprising the Virtual Chassis or VCF.

Virtual Chassis Fabric

Specify the role to be performed by each switch. In a VCF, two spine devices are configured into the Routing Engine role and the remaining spine devices and all leaf devices are configured into the `line-card` role. The role must also be associated with the member's serial number (see *serial-number*).

EX Series (except EX8200 Virtual Chassis) and QFX Series Virtual Chassis

Specify the role to be performed by each member switch. Associate the role with the member's serial number (see *serial-number*).

When you use a preprovisioned configuration, you cannot modify the primary-role priority or member ID of member switches through the user interfaces. The primary-role priority value is generated by the software, based on the assigned role:

- A member configured as `routing-engine` is assigned the primary-role priority 129.
- A member configured as `line-card` is assigned the primary-role priority 0.
- A member listed in the preprovisioned configuration without an explicitly specified role is assigned the `line-card` role and primary-role priority 0 by default.

The configured role specifications are permanent. If both `routing-engine` members fail, a `line-card` member cannot take over as primary of the Virtual Chassis configuration. You must delete the preprovisioned configuration to change the specified roles in a Virtual Chassis.

Explicitly configure two members as `routing-engine` and configure additional switches as members of the preprovisioned Virtual Chassis by specifying only their serial numbers. If you do not explicitly configure the role of the additional members, they function in a `line-card` role by default. In that case, a member that is functioning in a `line-card` role can take over primary role if the members functioning as primary and backup (`routing-engine` role) both fail.

EX8200 Virtual Chassis

Specify the role to be performed by each XRE200 External Routing Engine and each EX8200 member switch. Associate the role with the member's serial number (see *serial-number*). An EX8200 Virtual Chassis cannot function when both external Routing Engines, which must be configured in the `routing-engine` role, have failed.

Options

- `line-card`—Enables the member to be eligible to function only in the linecard role. Any member of the Virtual Chassis or VCF configuration other than the primary or backup functions in the linecard role and runs only a subset of Junos OS for EX Series switches. A member functioning in the linecard role does not run the control protocols or the chassis management processes.

  A Virtual Chassis must have at least three members for one member to function in the linecard role.

In an EX8200 Virtual Chassis configuration, all member switches must be in the linecard role.

- `routing-engine`—Enables the member to function as a primary or backup Routing Engine of the Virtual Chassis or VCF configuration. The primary manages all members and runs the chassis management processes and control protocols. The backup synchronizes with the primary in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the primary is unavailable.

  In a Virtual Chassis composed of EX Series switches (except EX8200 switches) or QFX Series switches, specify two and only two members in the `routing-engine` role. The software determines which of the two members assigned the `routing-engine` role functions as primary, based on the primary-role election algorithm. See Understanding How the Primary in a Virtual Chassis Is Elected. In these Virtual Chassis, the `routing-engine` role is associated with a switch. For a mixed Virtual Chassis or VCF, see Understanding Mixed EX Series and QFX Series Virtual Chassis or *Understanding Mixed Virtual Chassis Fabric* for specific recommendations or requirements for assigning the Routing Engine role based on the types of switches comprising the Virtual Chassis or VCF. The remaining switches are configured into the linecard role.

  In an EX8200 Virtual Chassis, all XRE200 External Routing Engines must be in the `routing-engine` role.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

Configuring an EX9200 Virtual Chassis

Configuring an EX4650 or a QFX Series Virtual Chassis

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

serial-number

IN THIS SECTION

Syntax

```
serial-number serial-number;
```

Hierarchy Level

```
[edit virtual-chassis member member-id]
```

Description

In a preprovisioned Virtual Chassis or Virtual Chassis Fabric (VCF), specify the serial number of each member switch to be included in the configuration. If you do not include the serial number within the configuration, the switch cannot be recognized as a member of a preprovisioned configuration. Serial number values are case-sensitive.

In an EX8200 Virtual Chassis configuration, specify the serial number of each XRE200 External Routing Engine and each EX8200 member switch to be included in the Virtual Chassis configuration. If you do not include the serial number within the Virtual Chassis configuration, the external Routing Engine or switch cannot be recognized as a member of the configuration.

Options

serial-number—Permanent serial number for the external Routing Engine or for the member switch.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

Configuring an EX9200 Virtual Chassis

Configuring an EX4650 or a QFX Series Virtual Chassis

traceoptions (Virtual Chassis)

IN THIS SECTION

Syntax

```
traceoptions  {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable | no-world-
readable>;
    flag flag <detail> <disable> <receive> <send>;
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Define tracing operations for the Virtual Chassis or VCF.

Each Virtual Chassis or VCF member is updated with configured tracing options, and the log file is stored locally on each device.

> **NOTE**: In a large-scale VCF, use system logging and tracing with discretion, and only for troubleshooting. These operations place an extra load on the primary Routing Engine device, which can impact VCF convergence time and stability. All tracing options should be disabled during normal VCF operation. To troubleshoot particular problems, selectively enable tracing options, and disable them again after collecting the desired information.

Default

Tracing operations are disabled.

Options

`detail`—(Optional) Generate detailed trace information for a flag.

> **NOTE**: Enable tracing at the detail level only while troubleshooting a particular issue, and disable it again for normal system operation.

`disable`—(Optional) Disable a flag.

`file` *`filename`*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *`number`*—(Optional) Maximum number of trace files. When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

`flag` *`flag`*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- `all`—All tracing operations.

> **TIP**: The `all` flag displays a subset of logs that are generally useful in debugging issues, and for more detailed information, you can use options `all detail`. However, to avoid significantly impacting VCF stability, use this high level of tracing only for troubleshooting for a short period of time, and not during normal VCF operation.

- `auto-configuration`—Trace Virtual Chassis ports (VCPs) that have been automatically configured.
- `csn`—Trace Virtual Chassis complete sequence number (CSN) packets.
- `error`—Trace Virtual Chassis errored packets.
- `hello`—Trace Virtual Chassis hello packets.
- `krt`—Trace Virtual Chassis KRT events.
- `lsp`—Trace Virtual Chassis link-state packets.
- `lsp-generation`—Trace Virtual Chassis link-state packet generation.
- `me`—Trace Virtual Chassis ME events.
- `normal`—Trace normal events.
- `packets`—Trace Virtual Chassis packets.
- `parse`—Trace reading of the configuration.
- `psn`—Trace partial sequence number (PSN) packets.
- `route`—Trace Virtual Chassis routing information.
- `spf`—Trace Virtual Chassis SPF events.
- `state`—Trace Virtual Chassis state transitions.
- `task`—Trace Virtual Chassis task operations.

`no-stamp`—(Optional) Do not place a timestamp on any trace file.

`no-world-readable`—(Optional) Restrict file access to the user who created the file.

`receive`—(Optional) Trace received packets.

`replace`—(Optional) Replace a trace file rather than appending information to it.

`send`—(Optional) Trace transmitted packets.

`size` *`size`*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *`trace-file`* reaches its maximum size, it is renamed *`trace-file`*`.0`, then *`trace-file`*`.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the `files` option.

- **Syntax:** *x*`k` to specify KB, *x*`m` to specify MB, or *x*`g` to specify GB
- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Option `detail` added in Junos OS Release 9.2 for EX Series switches.

RELATED DOCUMENTATION

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis
Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member
Verifying That Virtual Chassis Ports Are Operational
Troubleshooting an EX Series Virtual Chassis
Troubleshooting Virtual Chassis Fabric

upgrade-group

IN THIS SECTION

Syntax

```
upgrade-group group-name {
    fpcs (slot-number | [list-of-slot-numbers]);
    member member-id {
        fpcs (slot-number | [list-of-slot-numbers]);
    }
}
```

Hierarchy Level

```
[edit chassis nssu]
```

Description

Assign a name to a line-card upgrade group being created for nonstop software upgrade (NSSU).

To reduce the time an NSSU takes, you can configure line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines; EX8200 Virtual Chassis; EX4650 Virtual Chassis; QFX3500, QFX3600, and QFX5100 Virtual Chassis; and QFX5100 Virtual Chassis Fabric (VCF).

NSSU upgrades the groups in the order that they appear in the configuration (in other words, in the order you configure them). If you do not define any line-card upgrade groups, NSSU upgrades line cards or members of a Virtual Chassis or VCF one at a time in ascending order by slot or member number.

Options

group-name	Name of the upgrade group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade

vc-port

IN THIS SECTION

Syntax

```
vc-port {
    lag-hash (packet-based | source-port-based);
}
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Enable hashing of link aggregation group (LAG) network traffic over a dedicated trunk port within a Virtual Chassis, and select how the traffic within the dedicated trunk port is hashed.

BEST PRACTICE: Do not configure this statement unless you have a compelling reason to configure it. Configuration of this statement is optional and is only useful in a few types of network setups.

Default

source-port-based

Options

lag-hash	Select how to direct LAG traffic through the dedicated trunk port using either of the following options:
	packet-based — Hashes all incoming LAG network traffic on the dedicated trunk port based on the packet.
	source-port-based — Hashes all incoming LAG network traffic on the dedicated trunk port based on the source.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

vcp-no-hold-time

IN THIS SECTION

Syntax

```
vcp-no-hold-time;
```

Hierarchy Level

```
[edit virtual-chassis]
```

Description

Disable the Virtual Chassis port (VCP) holddown timer for all VCPs in the Virtual Chassis or Virtual Chassis Fabric (VCF).

The VCP holddown timer is an internal mechanism that delays a Virtual Chassis reconvergence for several seconds when a VCP becomes inactive. The purpose of this delay is to provide the VCP time to return online without having to reconverge the Virtual Chassis to adjust to the inactive VCP. All traffic to the VCP is dropped while the VCP is inactive. If the VCP remains down for a time that exceeds the VCP holddown timer, a Virtual Chassis reconvergence occurs.

This statement disables the holddown timer only in an EX4300 Virtual Chassis or a mixed Virtual Chassis that contains EX4300 switches in releases in the Junos OS 13.2X50 release train. In releases after that, the `vcp-no-hold-time` option is no longer needed and has no effect because the holddown timer is replaced by a planned PFE restart for actions that affect Virtual Chassis reconvergence. Switches and releases that don't support the holddown timer might allow you to configure this statement, but the configuration posts a warning message saying the statement has no effect. The option will be deprecated in an upcoming release and will no longer appear in the CLI.

When this statement is enabled, the VCP holddown timer is disabled and the Virtual Chassis reconvergence occurs when a VCP becomes inactive. The period of time where traffic is dropped waiting for the VCP to return online is avoided.

We recommend enabling this statement after a Virtual Chassis is operational. We recommend disabling this statement when you are adding or removing member switches from your Virtual Chassis.

The VCP holddown timer cannot be viewed and is not user-configurable. You can only control whether the VCP holddown timer is enabled or disabled by configuring this statement.

NOTE: In an EX4300 Virtual Chassis running a Junos OS 13.2X50 release, you should enable the `vcp-no-hold-time` statement before performing a software upgrade using NSSU. If you do not enable the `vcp-no-hold-time` statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see "Understanding Split and Merge in a Virtual Chassis" on page 71.

Default

The VCP holddown timer is enabled by default on all devices that support this statement.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

Understanding EX Series Virtual Chassis | 9
Understanding QFX Series Virtual Chassis | 25
Understanding Virtual Chassis Components | 35

virtual-chassis

IN THIS SECTION

Syntax

```
virtual-chassis {
    aliases {
        serial-number serial-number {
            alias-name alias-name;
        }
    }
    auto-conversion;
    auto-provisioned;
```

```
    auto-sw-update {
        (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
        package-name package-name;
    }
    fast-failover (ge | vcp disable | xe);
    graceful-restart {
        disable;
    }
    id id;
    mac-persistence-timer (minutes | disable);
    member member-id {
        fabric-tree-root;
        location location;
        mastership-priority number;
        no-management-vlan;
        serial-number;
        role;
    }
    no-auto-conversion;
    no-split-detection;
    preprovisioned;
    traceoptions (Virtual Chassis) {
        file filename <files number> <size size> <world-readable | no-world-readable> <match 
regex>;
        flag flag ;
    }
    vc-port {
        lag-hash (packet-based | source-port-based);
    }
    vcp-no-hold-time;
    vcp-snmp-statistics;
}
```

Hierarchy Level

```
[edit]
```

Description

Configure a Virtual Chassis or a Virtual Chassis Fabric (VCF).

The remaining statements are explained separately. See CLI Explorer.

Default

A standalone EX Series switch is a Virtual Chassis by default. It has a default member ID of 0, a default primary-role priority of 128, and a default role as primary.

A QFX Series device configured in standalone mode is a Virtual Chassis by default. It has a default member ID of 0, a default primary-role priority of 128, and a default role as primary.

A standalone XRE200 External Routing Engine or EX8200 switch is not part of an EX8200 Virtual Chassis until a Virtual Chassis configuration is set up.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 13.2X50-D15.

RELATED DOCUMENTATION

Autoprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric

Adding a Device to a Virtual Chassis Fabric

Configuring an EX4650 or a QFX Series Virtual Chassis

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

Configuring an EX9200 Virtual Chassis

6

CHAPTER

Operational Commands

clear virtual-chassis vc-port statistics

IN THIS SECTION

Syntax

```
clear virtual-chassis vc-port statistics
<all-members>
<interface-name>
<local>
<member member-id>
```

Description

Clear—reset to zero (0)—the traffic statistics counters on Virtual Chassis ports (VCPs).

Options

none	Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.
all-members	(Optional) Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.

interface-name	(Optional) Clear traffic statistics for the specified VCP.
local	(Optional) Clear traffic statistics for VCPs from the switch or external Routing Engine on which this command is entered.
member *member-id*	(Optional) Clear traffic statistics for VCPs from the specified member of a Virtual Chassis or VCF.

Required Privilege Level

clear

Sample Output

clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis)

```
user@switch> clear virtual-chassis vc-port statistics
fpc0:
--------------------------------------------------------------------------
Statistics cleared
```

clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis)

```
user@external-routing-engine> clear virtual-chassis vc-port statistics
member0:
--------------------------------------------------------------------------
Statistics cleared

member1:
--------------------------------------------------------------------------
Statistics cleared

member8:
--------------------------------------------------------------------------
Statistics cleared
```

```
member9:
--------------------------------------------------------------------------
Statistics cleared
```

clear virtual-chassis vc-port statistics member 3

```
user@switch> clear virtual-chassis vc-port statistics member 3
Cleared statistics on member 3
```

Release Information

Command introduced in Junos OS Release 9.0.

The options **all-members** and **local** were added in Junos OS Release 9.3 for EX Series switches.

RELATED DOCUMENTATION

show virtual-chassis vc-port statistics

show virtual-chassis vc-port

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

request session member

IN THIS SECTION

Syntax

```
request session member member-id
```

Description

Start a session with the specified member of a Virtual Chassis or a VCF.

Options

member-id Member ID for the specific member of the Virtual Chassis or VCF.

Required Privilege Level

maintenance

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

member

Understanding Virtual Chassis Components

request system software nonstop-upgrade

IN THIS SECTION

Syntax

```
request system software nonstop-upgrade (package-name | set [package-name package-name])
<force-host>
<no-copy>
<no-old-master-upgrade>
<reboot >
<unlink>
```

Description

Perform a nonstop software upgrade (NSSU) on a switch with redundant Routing Engines or on a Virtual Chassis or Virtual Chassis Fabric (VCF). The behavior of this command depends on the type of switch, Virtual Chassis, or VCF where you run it, as follows:

- When you run this command on any of the following Virtual Chassis or VCF configurations, NSSU upgrades all members of the Virtual Chassis:
    - EX3300, EX3400, EX4200, EX4300,EX4400, EX4500, EX4550, EX4600, or EX4650-48Y Virtual Chassis

  - Mixed Virtual Chassis composed of any combination of EX4200, EX4500, and EX4550 switches, or EX4300 and EX4600 switches
  - QFX3500 and QFX3600 Virtual Chassis
  - QFX5100 Virtual Chassis
  - QFX5120-48Y, QFX5120-48T or QFX5120-32C Virtual Chassis
  - Fixed configuration of switches in a VCF (QFX3500/QFX3600 and QFX5100 switches)
  - Mixed VCF composed of any combination of QFX3500/QFX3600, QFX5100, and EX4300 switches

  The original Virtual Chassis or VCF backup becomes the primary. The new primary automatically upgrades and reboots the original primary, which then rejoins the Virtual Chassis or VCF as the backup.

- When you run this command on an EX6200 or EX8200 switch, NSSU upgrades both the backup and primary Routing Engines. The original backup Routing Engine becomes the new primary at the end of the upgrade.
  - On an EX6200 switch, NSSU automatically reboots the original primary Routing Engine.
  - On an EX8200 switch, NSSU does not automatically reboot the original primary Routing Engine unless you specify the `reboot` option.
- When you run this command on an EX8200 Virtual Chassis, NSSU upgrades all primary and backup Routing Engines in the Virtual Chassis, including the external Routing Engines. The original backup Routing Engines become the new primary Routing Engines. NSSU does not automatically reboot the original primary Routing Engines unless you specify the `reboot` option.

This command has the following requirements:

- All Virtual Chassis members, VCF members, and all Routing Engines must be running the same Junos OS release.
- You must enable Graceful Routing Engine switchover (GRES)..
- You must enable Nonstop active routing (NSR).

> **NOTE**: Although not required, we recommend you also enable nonstop bridging (NSB). NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover during NSSU. See Configuring Nonstop Bridging on EX Series Switches (CLI Procedure).

- You must run the command from the primary Routing Engine on a standalone switch or from the primary on a Virtual Chassis.
- For minimal traffic disruption, you must define link aggregation groups (LAGs) such that the member links reside on different Virtual Chassis or VCF member switches (or on different line cards for EX6200 and EX8200 switches and EX8200 Virtual Chassis).
- For all Virtual Chassis (except EX8200 Virtual Chassis):
  - The Virtual Chassis members must be connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
  - The Virtual Chassis primary and backup must be adjacent to each other in the ring topology. With adjacent placement, the primary and backup are always in sync while the switches in line-card roles are rebooting.
  - The Virtual Chassis must be preprovisioned so the line-card role is explicitly assigned to member switches acting in a line-card role. During an NSSU, the primary and backup member switches must maintain their Routing Engine roles (although the primary role switches to the backup), and the remaining switches must maintain their line-card roles.
  - In a two-member Virtual Chassis, you must configure `no-split-detection` so the Virtual Chassis doesn't split during NSSU.
- For Virtual Chassis Fabric:
  - You can only have two members preprovisioned in the Routing Engine role. If more than two Routing Engines are configured, NSSU issues a warning message and the NSSU process stops.
  - The VCF members should be connected in a spine and leaf topology. A spine and leaf topology prevents the VCF from splitting during NSSU. Each leaf device must be connected to both spine devices.
  - The VCF must be preprovisioned so that the line-card role has been explicitly assigned to member switches acting in a line-card role, and likewise the Routing Engine role has been explicitly assigned to the member switches acting in a Routing Engine role. During an NSSU, the primary and backup member switches must maintain their Routing Engine roles (although the primary role switches to the backup), and the remaining switches must maintain their line-card roles.
  - You must configure `no-split-detection` in a two-member VCF so the VCF does not split during NSSU.

Options

package-name

Location of the software package or bundle to be installed. For example:

- `/var/tmp/`*`package-name`*—For a software package or bundle installed from a local directory on the switch.
- *`protocol`*`://`*`hostname`*`/`*`pathname`*`/`*`package-name`*—For a software package or bundle downloaded and installed from a remote location. Replace *`protocol`* with one of the following:
  - `ftp`—File Transfer Protocol.
    Use `ftp://`*`hostname`*`/`*`pathname`*`/`*`package-name`*.
    To specify authentication credentials, use `ftp://<`*`username`*`>:<`*`password`*`>@`*`hostname`*`/`*`pathname`*`/`*`package-name`*.
    To have the system prompt you for the password, specify `prompt` in place of the password.
    The command displays an error message If a password is required and you do not specify the password or `prompt`.
  - `http`—Hypertext Transfer Protocol.
    Use `http://`*`hostname`*`/`*`pathname`*`/`*`package-name`*.
    To specify authentication credentials, use `http://<`*`username`*`>:<`*`password`*`>@`*`hostname`*`/`*`pathname`*`/`*`package-name`*.
    The command prompts you for a password if one is required and you didn't include it.
  - `scp`—Secure copy (available only for Canada and U.S. version).
    Use `scp://`*`hostname`*`/`*`pathname`*`/`*`package-name`*.
    To specify authentication credentials, use `scp://<`*`username`*`>:<`*`password`*`>@`*`hostname`*`/`*`pathname`*`/`*`package-name`*.

> **NOTE**: The *`pathname`* in the protocol is the relative path to the user home directory on the remote system and not the root directory.

set [*package-name package-name*]

(Mixed Virtual Chassis only) Locations of the different installation packages required by the different types of member switches. These packages must be for the same Junos OS release. See this command's *`package-name`* option for information about how to specify the installation packages.

force-host

(Optional) Force adding the host software package or bundle (and ignore warnings) on EX4650, QFX5100, or QFX5120 devices.

no-copy (Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

no-old-master-upgrade (Optional) (EX8200 switches only) Upgrade the backup Routing Engine only. After the upgrade completes, the original primary Routing Engine becomes the backup Routing Engine and continues running the previous software version.

reboot (Optional) (EX8200 switches and EX8200 Virtual Chassis only) When you include the `reboot` option, NSSU automatically reboots the original primary (new backup) Routing Engine after being upgraded to the new software. When you omit the `reboot` option, you must manually reboot the original primary (new backup) Routing Engine using the `request system reboot` command.

> **NOTE**: If you do not use the `reboot` option on an EX8200 Virtual Chassis, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module to manually reboot the backup Routing Engines.

unlink (Optional) Remove the software package after a successful upgrade.

Required Privilege Level

maintenance

Output Fields

This command reports feedback on the status of the request. Some functions are shared between NSSU and the in-service software upgrade (ISSU) feature, so you might see what appear to be ISSU messages as well as NSSU messages in the output from this command.

Sample Output

request system software nonstop-upgrade (EX4200 Virtual Chassis)

```
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-4200-12.1R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Installing image on other FPC's along with the backup

Checking pending install on fpc1
Pushing bundle to fpc1
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc1

Checking pending install on fpc2
Pushing bundle to fpc2
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc2

Checking pending install on fpc3
Pushing bundle to fpc3
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc3

Checking pending install on fpc4
Pushing bundle to fpc4
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc4

Checking pending install on fpc5
Pushing bundle to fpc5
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc5
```

```
Checking pending install on fpc6
Pushing bundle to fpc6
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc6

Checking pending install on fpc7
Pushing bundle to fpc7
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Completed install on fpc7
Backup upgrade done
Rebooting Backup RE

Rebooting fpc1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE

*** FINAL System shutdown message from root@switch ***
```

```
System going down IMMEDIATELY



Shutdown NOW!
[pid 9336]
```

request system software nonstop-upgrade (EX6200 Switch)

```


{master}
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-6200-12.2R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re0
NOTICE: Validating configuration against jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re0
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online
```

```
  FPC 5          Online
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
  FPC 8          Online (ISSU)
  FPC 9          Online (ISSU)
Going to install image on master
NOTICE: Validating configuration against jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE
Trying to relinquish mastership before rebooting...
Resolving mastership...
Complete. The other routing engine becomes the master.

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY
```

request system software nonstop-upgrade reboot (EX8200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
```

```
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 2          Offline                 Offlined by CLI command
  FPC 3          Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
[pid 2635]


*** FINAL System shutdown message from user@switch ***
System going down IMMEDIATELY
```

request system software nonstop-upgrade no-old-master-upgrade (EX8200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade no-old-master-upgrade
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
```

```
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Offline                 Offlined by CLI command
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE
```

request system software nonstop-upgrade reboot (EX8200 Virtual Chassis)

```
{master:9}
user@external-routing-engine> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing LCC Backup REs
ISSU: Preparing Backup RE
Pushing bundle /var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz to member8
--------------------------------------------------------------------------
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
VC Backup upgrade done
Rebooting VC Backup RE

Rebooting member8
```

```
ISSU: Backup RE Prepare Done
Waiting for VC Backup RE reboot
Pushing bundle to member0-backup
Pushing bundle to member1-backup
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately

Rebooting member0-backup
Rebooting LCC [member0-backup]

Rebooting member1-backup
Rebooting LCC [member1-backup]
ISSU: LCC Backup REs Prepare Done
GRES operational
Initiating Chassis Nonstop-Software-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking Nonstop-Upgrade status
member0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 5          Online (ISSU)

member1:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Offline                 Offlined due to config
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 7          Online (ISSU)
```

```
member0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 5          Online (ISSU)

member1:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online (ISSU)
  FPC 1          Offline                 Offlined due to config
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 7          Online (ISSU)
ISSU: Upgrading Old Master RE
Pushing bundle /var/tmp/incoming-package-8200.tgz to member0-master
Pushing bundle /var/tmp/incoming-package-8200.tgz to member1-master

ISSU: RE switchover Done
WARNING: A reboot is required to install the software
WARNING:     Use the 'request system reboot' command immediately
Rebooting ...
shutdown: [pid 2188]
Shutdown NOW!
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!

*** FINAL System shutdown message from root@ ***
System going down IMMEDIATELY
```

Release Information

Command introduced in Junos OS Release 10.4.

Option `set [`*package-name package-name*`]` added in Junos OS Release 12.1 for EX Series switches.

RELATED DOCUMENTATION

show chassis nonstop-upgrade

Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)

Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade

Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade

request virtual-chassis mode

IN THIS SECTION

Syntax

```
request virtual-chassis mode
    fabric
    mixed <ieee-clause-82>
    network-port
    <disable>
    <reboot>
    <all-members> | <local> | <member member-id>
```

Description

Enable special device modes when operating as part of a Virtual Chassis or Virtual Chassis Fabric (VCF). Supported device modes include:

- Fabric mode—Enable the device to operate as part of a VCF. You must set `fabric` mode on all of the devices in a VCF.
- Mixed mode—Enable the device to operate in a mixed Virtual Chassis or VCF.
- (EX4400 switches only) Network port mode—Enable the default Virtual Chassis ports (VCPs) to operate as network ports instead of as VCPs.

You must reboot the standalone device, Virtual Chassis or VCF for this command to take effect, and the setting persists across subsequent reboot operations. You can use this command's `reboot` option to automatically reboot the device when setting or disabling a device mode.

> **NOTE**: For devices in a mixed VCF, you can set both mixed and fabric modes in the same command and then reboot the device once for both settings to take effect.

Mixed Mode Requirements and Best Practices

You must set `mixed` mode on all devices in a Virtual Chassis or a VCF when the devices have differences in how they interoperate. See Understanding Mixed EX Series and QFX Series Virtual Chassis or *Understanding Mixed Virtual Chassis Fabric* for details on which devices can be combined to form a mixed Virtual Chassis or VCF.

Some combinations of different types of switches form non-mixed Virtual Chassis or VCFs because the devices can run the same software images. You *don't* need to configure mixed mode if the only devices in your Virtual Chassis are the following combinations of switches:

- Only EX4500 and EX4550 switches in an EX Series Virtual Chassis.
- EX2300 multigigabit switches (EX2300-24MP and EX2300-48MP models) with any other EX2300 model switches in an EX2300 Virtual Chassis.
- EX4400 multigigabit switches (EX4400-24MP and EX4400-48MP models) with any other EX4400 model switches in an EX4400 Virtual Chassis.
- Only QFX3500 and QFX3600 switches in a QFX Series Virtual Chassis.
- QFX5110 and QFX5100 switches in a QFX5110 Virtual Chassis or QFX5110 VCF.

However, EX4300 switches that are not multigigabit models combined with EX4300 multigigabit model (EX4300-48MP) switches form a mixed EX4300 Virtual Chassis. In this case, you must set mixed mode on all member switches. In addition, when you set mixed mode on the EX4300 non-multigigabit model member switches, you must include the `ieee-clause-82` option with the mode command. This option sets a special port mode (IEEE Clause 82) on the Virtual Chassis ports (VCPs) that enables them to communicate when interconnected with VCPs on EX4300 multigigabit switches.

NOTE: If you remove an EX4300 member switch from a mixed EX4300 Virtual Chassis with multigigabit model members, you must disable `ieee-clause-82` port mode on the switch if you want to reconfigure it as a standalone switch or use it in any other type of mixed Virtual Chassis or any non-mixed Virtual Chassis. Otherwise, the VCPs will not connect with other members in the new Virtual Chassis. To do this, use the `request virtual-chassis mode mixed ieee-clause-82 disable` command and then reboot the switch. To add the switch to any other type of mixed Virtual Chassis, enable mixed mode again without this option and reboot the switch again for the change to take effect.

Best practices when setting `mixed` mode include:

- Don't enable mixed mode on a standalone device or for a member switch that is in a non-mixed Virtual Chassis or VCF. Mixed mode reduces the maximum scaling numbers for some features on the standalone switch, Virtual Chassis, or VCF to accommodate different types of devices that can be interconnected.

- To avoid potential traffic disruptions and configuration issues, on each device we recommend setting mixed mode for a mixed Virtual Chassis or VCF before cabling them into a mixed Virtual Chassis or VCF. You must then reboot the device, Virtual Chassis or VCF for the mode change to take effect.

  NOTE: The same best practice applies to setting fabric mode for a VCF—set this mode on the member devices before interconnecting them.

- You can change mode settings manually after a device has been added to a Virtual Chassis or VCF if you determine the settings are not the same on all the devices.

  However, if you set only some of the devices in a mixed Virtual Chassis or VCF to mixed mode, the mixed Virtual Chassis or VCF might not form. If you experience this issue, enter the `request virtual-chassis mode mixed all-members` command to force setting mixed mode on all devices in the Virtual Chassis or VCF. (Include the `fabric` mode option as well for a VCF.) You can then reboot the entire Virtual Chassis or VCF, or only the devices where the mode changed. The Virtual Chassis or VCF forms when the rebooted devices are online again.

NOTE: On switches running Junos OS Release 11.4 or later, if you don't set any of the switches in a mixed EX4200 and EX4500 Virtual Chassis to mixed mode, a mixed EX4200 and EX4500 Virtual Chassis forms where one of the switches assumes the primary role, and all other switches take on the linecard role. If you see this behavior, enter the `request virtual-chassis mode mixed all-members` command to set all members to mixed mode and reboot the Virtual Chassis.

Network Port Mode (EX4400 Switches Only)

If you are not using an EX4400 switch as a member switch in a Virtual Chassis, you can use this command with the `network-port` option to enable network port mode, which converts the default VCPs on the switch into network ports. By default, the two 100-Gigabit Ethernet QSFP28 ports on the rear panel are set as two logical 50-Gbps VCP interfaces each, forming a total of four logical 50-Gbps VCPs. (The default VCPs are in PIC slot 1, so the VCP interfaces are named vcp-255/1/0 through vcp-255/1/3.) If you convert these ports into network port mode, they become two 100 Gigabit Ethernet network ports, or you can channelize each of them into four 25-Gigabit Ethernet or four 10-Gigabit Ethernet network ports.

You can't use one of the rear panel 100-Gigabit Ethernet ports as a VCP and the other as a network port. When you use this command to enable network port mode, the mode affects both 100-Gbps ports together.

You can use the `show virtual-chassis mode` command to see the current port mode and what the port mode will be the next time the switch is rebooted. The two values might be different if you changed the mode but have not rebooted the switch yet.

To disable network port mode and return these ports to their default settings as VCPs, use the `network-port` and `disable` options with the `request virtual-chassis mode` command. You must reboot the switch for network port mode changes to take effect, so you can include the reboot option in the same command. For example:

```
request virtual-chassis mode network-port disable reboot
```

You can use the `request virtual-chassis vc-port delete` command to disable one of these ports as a VCP, but that action doesn't change the port mode to network port mode. You must enter the `request virtual-chassis mode network-port` command (and reboot the switch) to use the port as a network port. Also, if you run the `request virtual-chassis vc-port delete interface` *`interface-name`* command to disable vcp-255/1/0, that command disables both logical ports 0 and 1 (vcp-255/1/0 and vcp-255/1/1). Similarly, specifying to disable vcp-255/1/2 disables both logical ports 2 and 3 (vcp-255/1/2 and vcp-255/1/3).

Options

all-members (Optional) Set the Virtual Chassis mode for all members of the Virtual Chassis or VCF.

disable Disable the specified Virtual Chassis mode setting if it was previously enabled. You must include at least one of the supported modes that you want to disable with this option: `fabric`, `mixed`, or `network-port`.

The `request virtual-chassis mode mixed disable` command does not also automatically disable the mixed mode `ieee-clause-82` port mode option. To disable `ieee-clause-82` port mode on the device, you must include that option in the disable command as follows:

```
request virtual-chassis mode mixed ieee-clause-82 disable
```

fabric Set the device into fabric mode so it can participate as a member device in a VCF.

local (Optional) Set the Virtual Chassis mode only on the member device where the command is issued.

member *member-id* (Optional) Set the Virtual Chassis mode on the specified member of the Virtual Chassis or VCF.

mixed <ieee-clause-82> Set the device into mixed mode so that the device can participate in a mixed Virtual Chassis or mixed VCF.

You must set mixed mode with the `ieee-clause-82` option on EX4300 switches that are not multigigabit models when mixing them with EX4300 multigigabit (EX4300-48MP) model switches in an EX4300 Virtual Chassis.

network-port (EX4400 switches only) Set an EX4400 switch into network port mode, which converts the default VCPs on the switch to network port mode so you can use them as network ports instead of VCPs.

reboot After applying the mode change specified by the other options, reboot the device automatically. You must reboot the device for any mode change to take effect.

Required Privilege Level

system-control

Sample Output

request virtual-chassis mode mixed

```
user@switch> request virtual-chassis mode mixed
```

Sample Output

request virtual-chassis mode fabric mixed reboot

```
user@switch> request virtual-chassis mode fabric mixed reboot
```

Sample Output

request virtual-chassis mode mixed (with IEEE Clause 82 mode for EX4300 switches mixed with EX4300 multigigabit switches in a Virtual Chassis)

```
user@switch> request virtual-chassis mode ieee-clause-82 mixed
```

Release Information

Command introduced in Junos OS Release 11.1.

`fabric` mode introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices in a Virtual Chassis Fabric (VCF).

`ieee-clause-82` mixed-mode option introduced for EX4300 switches mixed with EX4300 multigigabit switches in a Virtual Chassis in Junos OS Release 18.2R1.

`network-port` mode introduced in Junos OS Release 21.1R1 for EX4400 switches in a Virtual Chassis.

RELATED DOCUMENTATION

Understanding Mixed EX Series and QFX Series Virtual Chassis

Understanding Mixed Virtual Chassis Fabric

Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis

Configuring an EX4650 or a QFX Series Virtual Chassis

Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

Verifying the Virtual Chassis Fabric Mode Settings

Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member

request virtual-chassis recycle

IN THIS SECTION

Syntax

```
request virtual-chassis recycle member-id member-id
```

Description

Make a previously used member ID available for reassignment.

When you remove a member switch from the Virtual Chassis configuration, the primary reserves that member ID. To make the member ID available for reassignment, you must use this command.

> **NOTE**: You must run this command from the Virtual Chassis member in the primary role.

Options

member-id ***member-id***	Specify the member ID that you want to make available for reassignment to a different member.

Required Privilege Level

system-control

Sample Output

request virtual-chassis recycle member-id 3

```
user@switch> request virtual-chassis recycle member-id 3
```

Sample Output

request virtual-chassis recycle member-id 1

```
user@external-routing-engine> request virtual-chassis recycle member-id 1
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

request virtual-chassis renumber

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

request virtual-chassis renumber

IN THIS SECTION

Syntax

```
request virtual-chassis renumber member-id old-member-id new-member-id new-member-id
```

Description

Renumber a member of a Virtual Chassis configuration.

NOTE: You must run this command from the Virtual Chassis member in the primary role.

Options

member-id *old-member-id*	Specify the ID of the member that you wish to renumber.
new-member-id *new-member-id*	Specify an unassigned member ID.

Required Privilege Level

system-control

Sample Output

request virtual-chassis renumber member-id 5 new-member-id 4

```
user@switch> request virtual-chassis renumber member-id 5 new-member-id 4
```

request virtual-chassis renumber member-id 1 new-member-id 0

```
user@external-routing-engine> request virtual-chassis renumber member-id 1 new-member-id 0
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

request virtual-chassis recycle

Removing or Replacing a Member Switch of a Virtual Chassis Configuration

request virtual-chassis vc-port

IN THIS SECTION

Syntax

```
request virtual-chassis vc-port (set | delete)
<fpc-slot fpc-slot> pic-slot pic-slot port port-number
interface interface-name
<member  | local>
```

member-id

Description

Set a port to operate as a Virtual Chassis port (VCP), or delete the VCP setting on a port. See Virtual Chassis Port Options for details on which ports you can set as VCPs on different switches. After setting a port as a VCP, you can't use the port for any other purpose unless you remove the VCP setting.

You can specify the port using the `interface` option or the `pic-slot` and `port` options.

If you don't include the `member` *`member-id`* option, this command defaults to setting the port as a VCP or deleting the VCP setting on the switch where you run the command. You can alternatively set the `local` option to ensure the command applies to the specified port on the local switch where you run it.

You might experience a temporary traffic disruption immediately after creating or deleting a user-configured VCP in an EX8200 Virtual Chassis.

Some switches have ports that are dedicated VCPs (you can only use them as VCPs and for no other purpose) or that are configured as VCPs by default (set as VCPs in the default factory configuration). You do not need to explicitly set those ports as VCPs to use them to interconnect the switch into a Virtual Chassis. See Virtual Chassis Port Options for details.

If you don't need a default-configured VCP to interconnect Virtual Chassis member switches, you can run this command with the `delete` option to remove the VCP setting and use the port as a regular uplink or network port. If you want to use the port as a VCP again later, then you need to run this command with the `set` option to reapply the VCP setting.

Special Behavior for VCPs on EX4400 Switches

On EX4400 switches, you don't use this command to convert ports on the switch into VCPs. EX4400 switches have default VCPs which are the only ports you can use as VCPs on the switch. These are the two 100-Gbps QSFP28 ports on the rear panel in PIC slot 1, which are set by default as four logical 50-Gbps VCP interfaces (vcp-255/1/0 through vcp-255/1/3). You can convert them into network ports by enabling network port mode on the switch (`request virtual-chassis mode network-port <reboot>`). In that case, to subsequently use the switch in a Virtual Chassis, you must disable network port mode (`request virtual-chassis mode network-port disable <reboot>`), which converts those ports back into VCPs. (Port mode changes apply to all VCPs on the switch,)

You can enter the `request virtual-chassis vc-port delete` command to disable these ports as VCPs, but, as mentioned above, that action doesn't convert them into network ports.

> **NOTE**: On EX4400 switches, if you use the `request virtual-chassis vc-port delete interface` *`interface-name`* command to disable logical VCP interface vcp-255/1/0, the switch disables both logical ports 0 and 1 (vcp-255/1/0 and vcp-255/1/1). Similarly, if you disable logical VCP interface vcp-255/1/2, that action disables both logical ports 2 and 3 (vcp-255/1/2 and vcp-255/1/3).

Options

set	Set a port as a VCP to convert an uplink or network port into a VCP.
delete	Delete the VCP setting on a port to convert a VCP back into an uplink or network port.
pic-slot *pic-slot*	Number of the PIC slot for the port on the switch.
port *port-number*	Number of the port that you want to enable or disable as a VCP.
interface *interface-name*	Interface name of the port that you want to enable or disable as a VCP. You can use this option to specify the port instead of using the `pic-slot` and `port` options.
member *member-id*	(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis or VCF.
local	(Optional) Enable or disable the specified VCP on the local switch where you run the command.

Required Privilege Level

system-control

Sample Output

request virtual-chassis vc-port set pic-slot 1 port 0

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port 0
```

To check the results of this command, use the *show virtual-chassis vc-port* command.

request virtual-chassis vc-port set pic-slot 1 port 1 member 3

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port 1 member 3
```

To check the results of this command, use the *show virtual-chassis vc-port* command.

request virtual-chassis vc-port delete pic-slot 1 port 1 member 3

```
user@switch> request virtual-chassis vc-port delete pic-slot 1 port 1 member 3
```

To check the results of this command, use the *show virtual-chassis vc-port* command.

Release Information

Command introduced in Junos OS Release 9.0.

Option `fpc-slot` introduced in Junos OS Release 10.4 for EX Series switches.

RELATED DOCUMENTATION

request virtual-chassis vc-port (Dedicated VCP)

show virtual-chassis vc-port

show virtual-chassis vc-port statistics

clear virtual-chassis vc-port statistics

Virtual Chassis Port (VCP) Interface Names in an EX8200 Virtual Chassis

Understanding Virtual Chassis Components

request virtual-chassis vc-port diagnostics optics

IN THIS SECTION

Syntax

```
request virtual-chassis vc-port diagnostics optics
```

Description

Run a digital optical monitoring (DOM) scan on the optical ports configured as Virtual Chassis ports (VCPs).

Enter the `show virtual-chassis vc-port diagnostics optics` command to view the results of the diagnostic scan.

On certain EX Series switches, you must enter the `request virtual-chassis vc-port diagnostics optics` command to run a diagnostic scan before you can gather the `show virtual-chassis vc-port diagnostics optics` output.

Required Privilege Level

system-control

Sample Output

request virtual-chassis vc-port diagnostics optics

```
user@switch> request virtual-chassis vc-port diagnostics optics
fpc0:
--------------------------------------------------------------------------
vc-port Diagnostics Optics Done
```

Release Information

Command introduced in Junos OS Release 13.2X50-D10.

RELATED DOCUMENTATION

show virtual-chassis vc-port diagnostics optics

show chassis nonstop-upgrade

IN THIS SECTION

Syntax

```
show chassis nonstop-upgrade
```

Description

(EX6200 switches, EX8200 switches, EX8200 Virtual Chassis, QFX3500 and QFX3600 Virtual Chassis, and Virtual Chassis Fabric only) Display the status of the line cards or Virtual Chassis members in the linecard role after the most recent nonstop software upgrade (NSSU). You must issue this command on the primary Routing Engine.

Required Privilege Level

view

Output Fields

Table 9 on page 286 lists the output fields for the `show chassis nonstop-upgrade` command. Output fields are listed in the approximate order in which they appear.

Table 9: show chassis nonstop-upgrade Output Fields

Field Name	Field Description
Item	Line card slot number.
Status	State of line card: • **Error**—Line card is in an error state. • **Offline**—Line card is powered down. • **Online**—Line card is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade (EX8200 Switch)

```
user@switch> show chassis nonstop-upgrade
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online
  FPC 3          Offline                 Offlined by CLI command
  FPC 4          Online
```

```
  FPC 5          Online
  FPC 6          Online
  FPC 7          Online
```

show chassis nonstop-upgrade (EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis nonstop-upgrade
member0:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online
  FPC 5          Online

member1:
--------------------------------------------------------------------------
  Item           Status                  Reason
  FPC 0          Online
  FPC 1          Offline                 Offlined due to config
  FPC 2          Online
  FPC 3          Online
  FPC 4          Online
  FPC 5          Online
  FPC 7          Online
```

show chassis nonstop-upgrade (Virtual Chassis Fabric)

```
 Item           Status                  Reason
  FPC 0          Online
  FPC 1          Online
  FPC 2          Online
  FPC 3          Online
  FPC 4          Online
  FPC 5          Online
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

request system software nonstop-upgrade

Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)

Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade

Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade

Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

show virtual-chassis

IN THIS SECTION

Syntax

```
show virtual-chassis <status>
```

Description

Display information about all members of the Virtual Chassis or VCF.

Options

none — Display information about all Virtual Chassis or VCF member devices.

status — Same output as for `show virtual-chassis` without any options.

Required Privilege Level

view

Output Fields

Table 10 on page 289 lists the output fields for the `show virtual-chassis` command. Output fields are listed in the approximate order in which they appear.

Table 10: show virtual-chassis Output Fields

Field Name	Field Description
Fabric ID	(VCF only) Assigned ID used to identify the VCF.
Fabric Mode	(VCF only) Mode of the VCF: Enabled, Disabled, or Mixed.
Preprovisioned Virtual Chassis or **Preprovisioned Virtual Chassis Fabric**	Virtual Chassis or VCF is configured using preprovisioning.

Table 10: show virtual-chassis Output Fields *(Continued)*

Field Name	Field Description
Virtual Chassis ID	Assigned ID that applies to the entire Virtual Chassis or VCF.
Virtual Chassis Mode	Mode of the Virtual Chassis or VCF. This field indicates support for the Virtual Chassis feature and, if a Virtual Chassis is configured, if it is a mixed or homogenous Virtual Chassis. Values can be: • `Enabled`—The platform supports the Virtual Chassis feature. If a Virtual Chassis is currently configured, this is a homogenous Virtual Chassis (all members are the same type of switch). • `Disabled`—The switch does not support the Virtual Chassis feature. **NOTE**: Switches that support the Virtual Chassis feature do not display this value. Even if a Virtual Chassis is not currently configured, those switches display `Enabled` in this field. • `Mixed`—The platform supports the Virtual Chassis feature, and is configured as a mixed mode Virtual Chassis (members consist of more than one type of switch).
Member ID	Assigned member ID and FPC: • On all EX Series Virtual Chassis except EX8200 Virtual Chassis, and on a VCF, the FPC number refers to the member ID assigned to the switch. • On EX8200 Virtual Chassis, member IDs are numbered 0 through 9. The FPC number indicates the slot number of the line card within the Virtual Chassis. The FPC number on member 0 is always 0 through 15. The FPC number on member 1 is always 16 through 31. The FPC number on member 2 is always 32 through 47; and so on for the members.

Table 10: show virtual-chassis Output Fields *(Continued)*

Field Name	Field Description
Status	For a nonprovisioned configuration: • `Prsnt` for a member that is currently connected to the Virtual Chassis or VCF configuration. • `NotPrsnt` for a member ID that has been assigned but is not currently connected. For a preprovisioned configuration: • `Prsnt` for a member that is specified in the preprovisioned configuration file and is currently connected to the Virtual Chassis or VCF. • `Unprvsnd` for a member that is interconnected with the Virtual Chassis or VCF configuration but is not specified in the preprovisioned configuration file.
Serial No	Serial number of the member device.
Alias-Name	The user-configured alias of the member device. The **Alias-Name** field appears only if an alias has been configured for at least one device in the Virtual Chassis or VCF. Aliases are configured using the `alias-name` statement in the `[edit virtual-chassis aliases serial-number` *`serial-number`*`]` hierarchy.
Model	Model number of the member device.
Mastership Priority or **Mstr prio**	Primary-role priority value of the member device.
Role	Role of the member device: primary, backup, or linecard. An asterisk (*) following the Role denotes the member device on which the `show virtual-chassis <status>` command was issued.

Table 10: show virtual-chassis Output Fields *(Continued)*

Field Name	Field Description
Mixed Mode	Mixed mode configuration status: • `Y` for a member device configured in mixed mode. • `N` for a member device not configured in mixed mode. • `NA` for a member device that cannot be configured in mixed mode.
Route Mode	The route mode of the member device: • `F` for fabric (VCF) mode • `VC` for Virtual Chassis mode
Location	Location of the member device. If this field is empty, the location field was not set for the device.
Neighbor List ID and **Interface**	Member ID of the neighbor member to which this member's Virtual Chassis port (VCP) is connected, and the VCP interface name.

Sample Output

show virtual-chassis (EX2300 multigigabit model—EX2300-24MP and EX2300-48MP—Virtual Chassis)

```
user@switch> show virtual-chassis
Virtual Chassis ID: cdef.789a.bc12
Virtual Chassis Mode: Enabled
                                           Mstr        Mixed Route  Neighbor List
Member ID  Status   Serial No  Model       prio Role    Mode  Mode   ID  Interface
0 (FPC 0)  Prsnt    ...        ex2300-24mp  128  Backup   N     VC    3  vcp-255/1/0
                                                                     1  vcp-255/1/2
1 (FPC 1)  Prsnt    ...        ex2300-48mp  128  Master*  N     VC    2  vcp-255/1/4
                                                                     0  vcp-255/1/1
```

```
2 (FPC 2)  Prsnt   ...        ex2300-24mp  128  Linecard N    VC   3  vcp-255/1/0
                                                                   1  vcp-255/1/1
3 (FPC 3)  Prsnt   ...        ex2300-48mp  128  Linecard N    VC   0  vcp-255/1/4
                                                                   2  vcp-255/1/5
```

show virtual-chassis (EX2300 and EX2300 multigigabit model (EX2300-24MP, EX2300-48MP) Virtual Chassis)

```
user@switch> show virtual-chassis status
Preprovisioned Virtual Chassis
Virtual Chassis ID: 9876.5432.abcd
Virtual Chassis Mode: Enabled
                                             Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No  Model         prio  Role      Mode Mode  ID  Interface
0 (FPC 0)  Prsnt    ...        ex2300-24mp   129   Backup      N  VC     1 vcp-255/1/3
                                                                         3 vcp-255/1/0
1 (FPC 1)  Prsnt    ...        ex2300-24p    129   Master*     N  VC     2 vcp-255/1/0
                                                                         0 vcp-255/1/1
2 (FPC 2)  Prsnt    ...        ex2300-24p      0   Linecard    N  VC     3 vcp-255/1/2
                                                                         1 vcp-255/1/3
3 (FPC 3)  Prsnt    ...        ex2300-48mp     0   Linecard    N  VC     2 vcp-255/1/5
                                                                         0 vcp-255/1/0
```

show virtual-chassis (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis
Virtual Chassis ID: 00ab.cdef.1234
Virtual Chassis Mode: Enabled
                                           Mastership      Mixed  Neighbor List
Member ID  Status   Serial No  Model       priority   Role  Mode  ID   Interface
0 (FPC 0)  Prsnt    ...        ex4200-24t    249     Master*  N   8    vcp-0
                                                                  1    vcp-1
1 (FPC 1)  Prsnt    ...        ex4200-24t    248     Backup   N   0    vcp-0
                                                                  2    vcp-1
2 (FPC 2)  Prsnt    ...        ex4200-48p    247    Linecard  N   1    vcp-0
                                                                  3    vcp-1
3 (FPC 3)  Prsnt    ...        ex4200-24t    246    Linecard  N   2    vcp-0
                                                                  4    vcp-1
```

```
4 (FPC 4)  Prsnt    ...        ex4200-48p   245   Linecard    N   3   vcp-0
                                                                   5   vcp-1
5 (FPC 5)  Prsnt    ...        ex4200-48t   244   Linecard    N   4   vcp-0
                                                                   6   vcp-1
6 (FPC 6)  Prsnt    ...        ex4200-48t   243   Linecard    N   5   vcp-0
                                                                   7   vcp-1
7 (FPC 7)  Prsnt    ...        ex4200-24f   242   Linecard    N   6   vcp-0
                                                                   8   vcp-1
8 (FPC 8)  Prsnt    ...        ex4200-24f   241   Linecard    N   7   vcp-0
                                                                   0   vcp-1
Member ID for next new member: 9 (FPC 9)
```

show virtual-chassis (Mixed EX4300 multigigabit model—EX4300-48MP—Virtual Chassis)

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis
Virtual Chassis ID: abcd.ef00.1234
Virtual Chassis Mode: Mixed
                                            Mstr         Mixed Route Neighbor List
Member ID   Status    Serial No  Model         prio  Role     Mode  Mode   ID Interface
0 (FPC 0)   Prsnt     ...        ex4300-48mp   129   Master*     Y   VC     1  vcp-255/1/0
                                                                            3  vcp-255/1/3
1 (FPC 1)   Prsnt     ...        ex4300-48mp   129   Backup      Y   VC     0  vcp-255/1/1
                                                                            2  vcp-255/1/3
2 (FPC 2)   Prsnt     ...        ex4300-48mp     0   Linecard    Y   VC     3  vcp-255/1/1
                                                                            1  vcp-255/1/3
3 (FPC 3)   Prsnt     ...        ex4300-48t      0   Linecard    Y   VC     0  vcp-255/1/0
                                                                            2  vcp-255/1/2
```

show virtual-chassis (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis
Virtual Chassis ID: cdcd.1212.efef
Virtual Chassis Mode: Enabled
                                        Mastership          Neighbor List
Member ID        Status Serial No  Model    priority  Role     ID Interface
0 (FPC 0-15)     Prsnt  ...        ex8216      0     Linecard  8  vcp-0/0
                                                               8  vcp-0/1
                                                               1  vcp-4/0/4
```

```
1 (FPC 16-31)   Prsnt  ...       ex8208      0    Linecard  8  vcp-0/0
                                                            0  vcp-3/0/4
8 (FPC 128-143) Prsnt  ...       ex-xre    128    Master    9  vcp-1/0
                                                            1  vcp-1/2
                                                            9  vcp-1/3
                                                            0  vcp-2/0
                                                            9  vcp-2/1
                                                            0  vcp-1/1
9 (FPC 144-159) Prsnt  ...       ex-xre    128    Backup*   8  vcp-1/0
                                                            8  vcp-1/2
                                                            8  vcp-1/3
                                                            8  vcp-1/3
```

show virtual-chassis (QFX5110 Virtual Chassis)

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis
Virtual Chassis ID: abab.1212.cdcd
Virtual Chassis Mode: Enabled
                                           Mstr         Mixed Route  Neighbor List
MemberID  Status Serial No  Model          prio  Role    Mode  Mode  ID Interface
0 (FPC 0) Prsnt  ...        qfx5110-32q    129   Backup     Y  VC    9 vcp-255/0/12
                                                                     1 vcp-255/0/16
1 (FPC 1) Prsn   ...        qfx5110-32q    129   Master*    Y  VC    2 vcp-255/0/16
                                                                     0 vcp-255/0/17
2 (FPC 2) Prsnt  ...        qfx5110-32q      0   Linecard   Y  VC    3 vcp-255/0/11
                                                                     1 vcp-255/0/14
3 (FPC 3) Prsnt  ...        qfx5110-32q      0   Linecard   Y  VC    4 vcp-255/0/14
                                                                     2 vcp-255/0/31
4 (FPC 4) Prsnt  ...        qfx5110-32q      0   Linecard   Y  VC    5 vcp-255/0/10
                                                                     3 vcp-255/0/5
5 (FPC 5) Prsnt  ...        qfx5100e-24q-2p  0   Linecard   Y  VC    6 vcp-255/0/12
                                                                     4 vcp-255/0/20
6 (FPC 6) Prsnt  ...        qfx5110-48s-4c   0   Linecard   Y  VC    7 vcp-255/0/10
                                                                     5 vcp-255/0/49
7 (FPC 7  Prsnt  ...        qfx5100e-48s-6q  0   Linecard   Y  VC    6 vcp-255/0/18
                                                                     8 vcp-255/0/31
8 (FPC 8) Prsnt  ...        qfx5110-48s-4c   0   Linecard   Y  VC    7 vcp-255/0/21
                                                                     9 vcp-255/0/49
```

```
9 (FPC 9) Prsnt  ...        qfx5110-32q      0   Linecard   Y  VC    8 vcp-255/0/13
                                                                     0 vcp-255/0/17
```

show virtual-chassis (QFX5200 Virtual Chassis)

```
user@switch> show virtual-chassis
Virtual Chassis ID: abab.1212.cdcd
Virtual Chassis Mode: Enabled
                 Serial                 Mstr           Mixed Route Neighbor List
Member ID Status   No    Model          prio  Role     Mode  Mode  ID Interface
0 (FPC 0) Prsnt   ...    qfx5200-32c-r  128  Master*   N     VC    1 vcp-255/0/19
                                                                   2 vcp-255/0/26
1 (FPC 1) Prsnt   ...    qfx5200-32c-r  128  Linecard  N     VC    0 vcp-255/0/20
                                                                   2 vcp-255/0/16
2 (FPC 2) Prsnt   ...    qfx5200-32c-32q 128  Backup   N     VC    0 vcp-255/0/15
                                                                   1 vcp-255/0/18
```

show virtual-chassis (QFX5100 Virtual Chassis Fabric)

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 0123.abcd.4567
Fabric Mode: Enabled
                   Serial               Mstr        Mixed Route Neighbor List
Member ID Status     No    Model        prio Role    Mode  Mode  ID Interface
0 (FPC 0) Prsnt    ...     qfx5100-48s   129  Master*  N    F     3  vcp-255/1/0
                                                                  2  vcp-255/1/1
                                                                  4  vcp-255/1/2
                                                                  4  vcp-255/1/3
1 (FPC 1) Prsnt    ...     qfx5100-48s   129  Backup   N    F     3  vcp-255/1/0
                                                                  2  vcp-255/1/1
                                                                  4  vcp-255/1/2
                                                                  4  vcp-255/1/3
2 (FPC 2) Prsnt    ...     qfx5100-48s     0  Linecard N    F     1  vcp-255/1/0
                                                                  0  vcp-255/1/1
3 (FPC 3) Prsnt    ...     qfx5100-48s     0  Linecard N    F     1  vcp-255/1/0
                                                                  0  vcp-255/1/1
4 (FPC 4) Prsnt    ...     qfx5100-48s     0  Linecard N    F     1  vcp-255/1/0
                                                                  0  vcp-255/1/1
```

Release Information

Command introduced in Junos OS Release 9.2.

`Fabric ID`, `Fabric Mode`, and `Route Mode` output fields introduced in Junos OS Release 13.2X51-D20.

`Alias-Name` output field introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

show virtual-chassis active-topology

show virtual-chassis protocol adjacency

show virtual-chassis vc-path

Understanding Mixed EX Series and QFX Series Virtual Chassis

Understanding Mixed Virtual Chassis Fabric

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

show virtual-chassis active-topology

IN THIS SECTION

Syntax

```
show virtual-chassis active-topology
<all-members | local | member member-id>
```

Description

Display the active topology of the Virtual Chassis or VCF with next-hop reachability information.

Options

none	Display the active topology of the member switch where you enter this command.
all-members	(Optional) Display the active topology of all members of the Virtual Chassis or VCF.
local	(Optional) Display the active topology of the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display the active topology of the specified member of the Virtual Chassis or VCF.

Required Privilege Level

view

Output Fields

Table 11 on page 299 lists the output fields for the `show virtual-chassis active-topology` command. Output fields are listed in the approximate order in which they appear.

Table 11: show virtual-chassis active-topology Output Fields

Field Name	Field Description
Destination ID	Specifies the member ID of the destination.
Next-hop	Specifies the member ID and Virtual Chassis port (VCP) of the next hop to which packets for the destination ID are forwarded. The next hop can be more than one device in a VCF.

Sample Output

show virtual-chassis active-topology (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis active-topology

  1                   1(vcp-1)



  2                   1(vcp-1)



  3                   1(vcp-1)



  4                   1(vcp-1)



  5                   8(vcp-0)  1(vcp-1)



  6                   8(vcp-0)
```

```
  7                    8(vcp-0)

  8                    8(vcp-0)
```

show virtual-chassis active-topology (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis active-topology

 member0:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  1                    1(vcp-4/0/4.32768)

  8                    8(vcp-0/0.32768)

  9                    8(vcp-0/0.32768)

member1:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-3/0/4.32768)

  8                    8(vcp-0/0.32768)

  9                    8(vcp-0/0.32768)

member8:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-1/1.32768)
```

```
  1                    1(vcp-1/2.32768)

  9                    9(vcp-2/1.32768)

member9:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    8(vcp-1/2.32768)

  1                    8(vcp-1/2.32768)

  8                    8(vcp-1/2.32768)
```

show virtual-chassis active-topology (Virtual Chassis Fabric)

```
user@device> show virtual-chassis active-topology
fpc0:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  1                    4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  2                    4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  3                    4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  4                    4(vcp-255/0/2.32768)

  5                    5(vcp-255/0/3.32768)

  6                    6(vcp-255/0/1.32768)

fpc1:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  2                    4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)
```

```
  3                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  4                     4(vcp-255/0/2.32768)

  5                     5(vcp-255/0/3.32768)

  6                     6(vcp-255/0/1.32768)

fpc2:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  1                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  3                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  4                     4(vcp-255/0/2.32768)

  5                     5(vcp-255/0/3.32768)

  6                     6(vcp-255/0/1.32768)

fpc3:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  1                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  2                     4(vcp-255/0/2.32768)  5(vcp-255/0/3.32768)  6(vcp-255/0/1.32768)

  4                     4(vcp-255/0/2.32768)

  5                     5(vcp-255/0/3.32768)

  6                     6(vcp-255/0/1.32768)

fpc4:
--------------------------------------------------------------------------
  Destination ID        Next-hop
```

```
  0                    0(vcp-255/0/48.32768)

  1                    1(vcp-255/0/49.32768)

  2                    2(vcp-255/0/50.32768)

  3                    3(vcp-255/0/51.32768)

  5                    3(vcp-255/0/51.32768)  2(vcp-255/0/50.32768)  0(vcp-255/0/48.32768)
1(vcp-255/0/49.32768)

  6                    3(vcp-255/0/51.32768)  2(vcp-255/0/50.32768)  0(vcp-255/0/48.32768)
1(vcp-255/0/49.32768)

fpc5:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-255/0/48.32768)

  1                    1(vcp-255/0/49.32768)

  2                    2(vcp-255/0/50.32768)

  3                    3(vcp-255/0/51.32768)

  4                    3(vcp-255/0/51.32768)  2(vcp-255/0/50.32768)  0(vcp-255/0/48.32768)
1(vcp-255/0/49.32768)

  6                    3(vcp-255/0/51.32768)  2(vcp-255/0/50.32768)  0(vcp-255/0/48.32768)
1(vcp-255/0/49.32768)

fpc6:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-255/0/0.32768)

  1                    1(vcp-255/0/1.32768)

  2                    2(vcp-255/0/2.32768)
```

```
  3                   3(vcp-255/0/3.32768)

  4                   3(vcp-255/0/3.32768)  2(vcp-255/0/2.32768)  0(vcp-255/0/0.32768)
1(vcp-255/0/1.32768)

  5                   3(vcp-255/0/3.32768)  2(vcp-255/0/2.32768)  0(vcp-255/0/0.32768)
1(vcp-255/0/1.32768)
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

show virtual-chassis device-topology

IN THIS SECTION

Syntax

```
show virtual-chassis device-topology
<all-members | local | member member-id>
```

Description

Display the device topology—the member and system IDs, the VCP numbers, and device status—for all hardware devices in the Virtual Chassis or VCF.

Options

none	Display the device topology for all members of the Virtual Chassis or VCF.
all-members	(Optional) Display the device topology for all members of the Virtual Chassis or VCF.
local	(Optional) Display the device topology for the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display the device topology for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 12 on page 306 lists the output fields for the `show virtual-chassis device-topology` command. Output fields are listed in the approximate order in which they appear.

Table 12: show virtual-chassis device-topology Output Fields

Field Name	Field Description
Member	Assigned member ID.
Device	Assigned device ID. For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.
Status	The status of the device within the Virtual Chassis or VCF. Outputs include: • **Prsnt**—Device is currently connected to and participating in the Virtual Chassis or VCF. • **NotPrsnt**—Device is assigned but is not currently connected.
System ID	System ID of the device. The system ID of the device is the device's MAC address.
Member (Neighbor List)	Assigned member ID of the neighbor device.
Device (Neighbor List)	Assigned device ID of the neighbor device. For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.
Interface (Neighbor List)	The interface connecting the device to the neighbor.

Sample Output

show virtual-chassis device-topology

```
user@switch> show virtual-chassis device-topology
member0:
--------------------------------------------------------------------------
```

```
                                              Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    0021.59f7.d000        8       8    vcp-0/0
                                               1       1    vcp-4/0/1
   1       1    Prsnt    0026.888d.6800        8       8    vcp-0/0
                                               9       9    vcp-0/1
                                               0       0    vcp-3/0/4
   8       8    Prsnt    0000.4a75.9b7c        9       9    vcp-1/0
                                               0       0    vcp-1/1
                                               1       1    vcp-1/2
   9       9    Prsnt    0000.73e9.9a57        8       8    vcp-1/0
                                               1       1    vcp-1/1

member1:
--------------------------------------------------------------------------
                                              Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    0021.59f7.d000        8       8    vcp-0/0
                                               1       1    vcp-4/0/1
   1       1    Prsnt    0026.888d.6800        8       8    vcp-0/0
                                               9       9    vcp-0/1
                                               0       0    vcp-3/0/4
   8       8    Prsnt    0000.4a75.9b7c        9       9    vcp-1/0
                                               0       0    vcp-1/1
                                               1       1    vcp-1/2
   9       9    Prsnt    0000.73e9.9a57        8       8    vcp-1/0
                                               1       1    vcp-1/1

member8:
--------------------------------------------------------------------------
                                              Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    0021.59f7.d000        8       8    vcp-0/0
                                               1       1    vcp-4/0/1
   1       1    Prsnt    0026.888d.6800        8       8    vcp-0/0
                                               9       9    vcp-0/1
                                               0       0    vcp-3/0/4
   8       8    Prsnt    0000.4a75.9b7c        9       9    vcp-1/0
                                               0       0    vcp-1/1
                                               1       1    vcp-1/2
   9       9    Prsnt    0000.73e9.9a57        8       8    vcp-1/0
                                               1       1    vcp-1/1
```

```
member9:
--------------------------------------------------------------------------
                                              Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    0021.59f7.d000       8       8    vcp-0/0
                                              1       1    vcp-4/0/1
   1       1    Prsnt    0026.888d.6800       8       8    vcp-0/0
                                              9       9    vcp-0/1
                                              0       0    vcp-3/0/4
   8       8    Prsnt    0000.4a75.9b7c       9       9    vcp-1/0
                                              0       0    vcp-1/1
                                              1       1    vcp-1/2
   9       9    Prsnt    0000.73e9.9a57       8       8    vcp-1/0
                                              1       1    vcp-1/1
```

show virtual-chassis device-topology (Virtual Chassis Fabric)

```
user@device> show virtual-chassis device-topology
fpc0:
--------------------------------------------------------------------------
                                              Neighbor List
Member  Device  Status   System ID         Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900       4       4    vcp-255/0/2
                                              5       5    vcp-255/0/3
                                              6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40       4       4    vcp-255/0/2
                                              5       5    vcp-255/0/3
                                              6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700       4       4    vcp-255/0/2
                                              5       5    vcp-255/0/3
                                              6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440       4       4    vcp-255/0/2
                                              5       5    vcp-255/0/3
                                              6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40       3       3    vcp-255/0/51
                                              2       2    vcp-255/0/50
                                              0       0    vcp-255/0/48
                                              1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0       3       3    vcp-255/0/51
                                              2       2    vcp-255/0/50
                                              1       1    vcp-255/0/49
```

```
                                                  0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00           3       3    vcp-255/0/3
                                                  2       2    vcp-255/0/2
                                                  0       0    vcp-255/0/0
                                                  1       1    vcp-255/0/1

fpc1:
--------------------------------------------------------------------------
                                                   Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900           4       4    vcp-255/0/2
                                                  5       5    vcp-255/0/3
                                                  6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40           4       4    vcp-255/0/2
                                                  5       5    vcp-255/0/3
                                                  6       6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700           4       4    vcp-255/0/2
                                                  5       5    vcp-255/0/3
                                                  6       6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440           4       4    vcp-255/0/2
                                                  5       5    vcp-255/0/3
                                                  6       6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40           3       3    vcp-255/0/51
                                                  2       2    vcp-255/0/50
                                                  0       0    vcp-255/0/48
                                                  1       1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0           3       3    vcp-255/0/51
                                                  2       2    vcp-255/0/50
                                                  1       1    vcp-255/0/49
                                                  0       0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00           3       3    vcp-255/0/3
                                                  2       2    vcp-255/0/2
                                                  0       0    vcp-255/0/0
                                                  1       1    vcp-255/0/1

fpc2:
--------------------------------------------------------------------------
                                                   Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900           4       4    vcp-255/0/2
                                                  5       5    vcp-255/0/3
                                                  6       6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40           4       4    vcp-255/0/2
```

```
                                              5         5    vcp-255/0/3
                                              6         6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700       4         4    vcp-255/0/2
                                              5         5    vcp-255/0/3
                                              6         6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440       4         4    vcp-255/0/2
                                              5         5    vcp-255/0/3
                                              6         6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40       3         3    vcp-255/0/51
                                              2         2    vcp-255/0/50
                                              0         0    vcp-255/0/48
                                              1         1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0       3         3    vcp-255/0/51
                                              2         2    vcp-255/0/50
                                              1         1    vcp-255/0/49
                                              0         0    vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00       3         3    vcp-255/0/3
                                              2         2    vcp-255/0/2
                                              0         0    vcp-255/0/0
                                              1         1    vcp-255/0/1

fpc3:
--------------------------------------------------------------------------
                                               Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900       4         4    vcp-255/0/2
                                              5         5    vcp-255/0/3
                                              6         6    vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40       4         4    vcp-255/0/2
                                              5         5    vcp-255/0/3
                                              6         6    vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700       4         4    vcp-255/0/2
                                              5         5    vcp-255/0/3
                                              6         6    vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440       4         4    vcp-255/0/2
                                              5         5    vcp-255/0/3
                                              6         6    vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40       3         3    vcp-255/0/51
                                              2         2    vcp-255/0/50
                                              0         0    vcp-255/0/48
                                              1         1    vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0       3         3    vcp-255/0/51
                                              2         2    vcp-255/0/50
```

```
                                                 1        1     vcp-255/0/49
                                                 0        0     vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00          3        3     vcp-255/0/3
                                                 2        2     vcp-255/0/2
                                                 0        0     vcp-255/0/0
                                                 1        1     vcp-255/0/1

fpc4:
--------------------------------------------------------------------------
                                                  Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900          4        4     vcp-255/0/2
                                                 5        5     vcp-255/0/3
                                                 6        6     vcp-255/0/1
   1       1    Prsnt    100e.7eb8.3a40          4        4     vcp-255/0/2
                                                 5        5     vcp-255/0/3
                                                 6        6     vcp-255/0/1
   2       2    Prsnt    100e.7eb5.d700          4        4     vcp-255/0/2
                                                 5        5     vcp-255/0/3
                                                 6        6     vcp-255/0/1
   3       3    Prsnt    100e.7eb5.c440          4        4     vcp-255/0/2
                                                 5        5     vcp-255/0/3
                                                 6        6     vcp-255/0/1
   4       4    Prsnt    100e.7eb5.7e40          3        3     vcp-255/0/51
                                                 2        2     vcp-255/0/50
                                                 0        0     vcp-255/0/48
                                                 1        1     vcp-255/0/49
   5       5    Prsnt    100e.7eb5.80c0          3        3     vcp-255/0/51
                                                 2        2     vcp-255/0/50
                                                 1        1     vcp-255/0/49
                                                 0        0     vcp-255/0/48
   6       6    Prsnt    100e.7eb6.3b00          3        3     vcp-255/0/3
                                                 2        2     vcp-255/0/2
                                                 0        0     vcp-255/0/0
                                                 1        1     vcp-255/0/1

fpc5:
--------------------------------------------------------------------------
                                                  Neighbor List
Member  Device  Status   System ID          Member  Device  Interface
   0       0    Prsnt    100e.7eb6.a900          4        4     vcp-255/0/2
                                                 5        5     vcp-255/0/3
                                                 6        6     vcp-255/0/1
```

```
   1        1     Prsnt     100e.7eb8.3a40        4        4     vcp-255/0/2
                                                  5        5     vcp-255/0/3
                                                  6        6     vcp-255/0/1
   2        2     Prsnt     100e.7eb5.d700        4        4     vcp-255/0/2
                                                  5        5     vcp-255/0/3
                                                  6        6     vcp-255/0/1
   3        3     Prsnt     100e.7eb5.c440        4        4     vcp-255/0/2
                                                  5        5     vcp-255/0/3
                                                  6        6     vcp-255/0/1
   4        4     Prsnt     100e.7eb5.7e40        3        3     vcp-255/0/51
                                                  2        2     vcp-255/0/50
                                                  0        0     vcp-255/0/48
                                                  1        1     vcp-255/0/49
   5        5     Prsnt     100e.7eb5.80c0        3        3     vcp-255/0/51
                                                  2        2     vcp-255/0/50
                                                  1        1     vcp-255/0/49
                                                  0        0     vcp-255/0/48
   6        6     Prsnt     100e.7eb6.3b00        3        3     vcp-255/0/3
                                                  2        2     vcp-255/0/2
                                                  0        0     vcp-255/0/0
                                                  1        1     vcp-255/0/1

fpc6:
--------------------------------------------------------------------------
                                                   Neighbor List
Member  Device  Status    System ID          Member  Device  Interface
   0        0     Prsnt     100e.7eb6.a900        4        4     vcp-255/0/2
                                                  5        5     vcp-255/0/3
                                                  6        6     vcp-255/0/1
   1        1     Prsnt     100e.7eb8.3a40        4        4     vcp-255/0/2
                                                  5        5     vcp-255/0/3
                                                  6        6     vcp-255/0/1
   2        2     Prsnt     100e.7eb5.d700        4        4     vcp-255/0/2
                                                  5        5     vcp-255/0/3
                                                  6        6     vcp-255/0/1
   3        3     Prsnt     100e.7eb5.c440        4        4     vcp-255/0/2
                                                  5        5     vcp-255/0/3
                                                  6        6     vcp-255/0/1
   4        4     Prsnt     100e.7eb5.7e40        3        3     vcp-255/0/51
                                                  2        2     vcp-255/0/50
                                                  0        0     vcp-255/0/48
                                                  1        1     vcp-255/0/49
   5        5     Prsnt     100e.7eb5.80c0        3        3     vcp-255/0/51
```

```
                                        2      2    vcp-255/0/50
                                        1      1    vcp-255/0/49
                                        0      0    vcp-255/0/48
  6       6    Prsnt    100e.7eb6.3b00  3      3    vcp-255/0/3
                                        2      2    vcp-255/0/2
                                        0      0    vcp-255/0/0
                                        1      1    vcp-255/0/1
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Virtual Chassis Port Link Aggregation

Understanding EX8200 Virtual Chassis Topologies

show virtual-chassis login

IN THIS SECTION

Syntax

```
show virtual-chassis login
```

Description

Supply the address of the host that logged into the Virtual Chassis or VCF, or identify the location of the member switch that redirected the current session to a different member switch.

You might need this information for tracing or troubleshooting purposes.

Required Privilege Level

view

Sample Output

show virtual-chassis login (Direct Login to the Primary Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host 192.0.2.3
```

show virtual-chassis login (Backup Console Session Redirected to the Primary Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host backup
```

Release Information

Command introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

request session member

Understanding Global Management of a Virtual Chassis

show virtual-chassis mode

IN THIS SECTION

Syntax

```
show virtual-chassis mode
<all-members | local | member member-id>
```

Description

Display the Virtual Chassis or Virtual Chassis Fabric (VCF) mixed mode status.

Options

none	Display the Virtual Chassis or VCF mixed mode status for the device where you enter the command.
all-members	(Optional) Display the Virtual Chassis or VCF mixed mode status for all member devices in the Virtual Chassis or VCF.
local	(Optional) Display the Virtual Chassis or VCF mixed mode status for the device where you enter the command.
member *member-id*	(Optional) Display the Virtual Chassis or VCF mixed mode status for the specified member device.

Required Privilege Level

view

Output Fields

Table 13 on page 316 lists the output fields for the `show virtual-chassis mode` command.

Table 13: show virtual-chassis mode Output Fields

Field Name	Field Description
Mixed Mode	Specifies the mixed mode status of the member switch. Mixed mode is either **Enabled** or **Disabled**.

Table 13: show virtual-chassis mode Output Fields *(Continued)*

Field Name	Field Description
Current mode	Shows the current mixed and fabric mode settings running on the member device or devices. (EX4400 switches only) Shows whether the switch has network port mode enabled to use the default Virtual Chassis ports (VCPs) as network ports instead of as VCPs. A device reboot is required to change the device mode. The **Current mode** and **Future mode after reboot** are different when the mode has been changed but the device has not been rebooted. Outputs include: • **Fabric with mixed devices**—Fabric mode and mixed mode are enabled. • **Fabric with similar devices**—Fabric mode is enabled and mixed mode is disabled. • **Virtual Chassis with mixed devices**—Fabric mode is disabled and mixed mode is enabled. • **Virtual Chassis with similar devices**—Fabric mode is disabled and mixed mode is disabled. • (EX4400 switches only) **network-port mode**—Network port mode is enabled, which means the default VCPs on the switch function as network ports rather than as VCPs.
Future mode after reboot	Shows the mixed and fabric mode settings that you have requested to be enabled on the member device or devices the next time the device is rebooted. A device reboot is required to change the fabric or mixed mode. The **Current mode** and **Future mode after reboot** are different only when the mode has been changed but the device has not been rebooted. Outputs include: • **Fabric with mixed devices**—Fabric mode and mixed mode are enabled. • **Fabric with similar devices**—Fabric mode is enabled and mixed mode is disabled. • **Virtual Chassis with mixed devices**—Fabric mode is disabled and mixed mode is enabled. • **Virtual Chassis with similar devices**—Fabric mode is disabled and mixed mode is disabled. • (EX4400 switches only) **network-port mode**—Network port mode is enabled so the default VCPs on the switch will function as network ports rather than as VCPs.

Sample Output

show virtual-chassis mode (EX4200 in a non-mixed Virtual Chassis)

```
user@switch> show virtual-chassis mode

fpc0:
--------------------------------------------------------------------------
Mixed Mode: Disabled
```

Sample Output

show virtual-chassis mode (EX4400 with network port mode enabled)

```
user@switch> show virtual-chassis mode

fpc0:
--------------------------------------------------------------------------
Current mode : network-port mode
Future mode after reboot: network-port mode
```

Sample Output

show virtual-chassis mode (QFX5100 in a non-mixed VCF)

```
user@switch> show virtual-chassis mode

fpc0:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc1:
--------------------------------------------------------------------------
```

```
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc2:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc3:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc4:
--------------------------------------------------------------------------
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices
```

Release Information

Command introduced in Junos OS Release 11.1.

Current mode and **Future mode after reboot** fields introduced in Junos OS Release 13.2X51-D20.

RELATED DOCUMENTATION

request virtual-chassis mode

Verifying the Virtual Chassis Fabric Mode Settings

Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)

show virtual-chassis protocol adjacency

IN THIS SECTION

Syntax

```
show virtual-chassis protocol adjacency
<brief | detail | extensive>
<all-members | local | member member-id>
<system-id>
```

Description

Display the Virtual Chassis Control Protocol (VCCP) adjacency statistics in the Virtual Chassis or VCF for all hardware devices.

Options

none	Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.

brief | detail | extensive	(Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` and `extensive` options provide identical displays.
all-members	(Optional) Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.
local	(Optional) Display VCCP adjacency statistics for the switch or external Routing Engine on which this command is entered.
member *member-id*	(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.
system-id	(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 14 on page 321 lists the output fields for the `show virtual-chassis protocol adjacency` command. Output fields are listed in the approximate order in which they appear.

Table 14: show virtual-chassis protocol adjacency Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the Virtual Chassis port (VCP) interface.	All levels
`System`	The MAC address of the device on the receiving side of the VCP link.	All levels

Table 14: show virtual-chassis protocol adjacency Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`State`	State of the link. Outputs include: • `Up`—The link is up. • `Down`—The link is down. • `New`—The link is new. • `One-way`—The link is transmitting traffic in one direction. • `Initializing`—The link is initializing. • `Rejected`—The link is rejected.	All levels
`Hold, Expires in`	Remaining holdtime of the adjacency.	All levels
`Priority`	Priority to become the designated intermediary system.	detail
`Up/Down Transitions`	Count of adjacency status transition changes from up to down or down to up.	detail
`Last transition`	Time of the last up/down transition.	detail

Sample Output

show virtual-chassis protocol adjacency

```
user@switch> show virtual-chassis protocol adjacency
member0:
--------------------------------------------------------------------------
Interface             System         State        Hold (secs)
vcp-0/0.32768         0000.4a75.9b7c Up                  57
vcp-0/1.32768         0000.4a75.9b7c Up                  59
vcp-4/0/1.32768       0026.888d.6800 Up                  57
```

```
member1:
--------------------------------------------------------------------------
Interface           System         State        Hold (secs)
vcp-0/0.32768       0000.4a75.9b7c Up                  58
vcp-0/1.32768       0000.73e9.9a57 Up                  59
vcp-3/0/4.32768     0021.59f7.d000 Up                  58

member8:
--------------------------------------------------------------------------
Interface           System         State        Hold (secs)
vcp-1/0.32768       0000.73e9.9a57 Up                  58
vcp-1/1.32768       0021.59f7.d000 Up                  58
vcp-1/2.32768       0026.888d.6800 Up                  59
vcp-2/0.32768       0021.59f7.d000 Up                  59

member9:
--------------------------------------------------------------------------
Interface           System         State        Hold (secs)
vcp-1/0.32768       0000.4a75.9b7c Up                  58
vcp-1/1.32768       0026.888d.6800 Up                  59
```

show virtual-chassis protocol adjacency detail

```
user@switch> show virtual-chassis protocol adjacency detail
member0:
--------------------------------------------------------------------------


0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago

0000.4a75.9b7c
  interface-name: vcp-0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago

0026.888d.6800
  interface-name: vcp-4/0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago

member1:
```

```
----------------------------------------------------------------------------

0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0000.73e9.9a57
  interface-name: vcp-0/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago

0021.59f7.d000
  interface-name: vcp-3/0/4.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago

member8:
----------------------------------------------------------------------------

0000.73e9.9a57
  interface-name: vcp-1/0.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0021.59f7.d000
  interface-name: vcp-1/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0026.888d.6800
  interface-name: vcp-1/2.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0021.59f7.d000
  interface-name: vcp-2/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

member9:
----------------------------------------------------------------------------

0000.4a75.9b7c
  interface-name: vcp-1/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago

0026.888d.6800
```

```
interface-name: vcp-1/1.32768, State: Up, Expires in 58 secs
Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Virtual Chassis Port Link Aggregation

Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

show virtual-chassis protocol database

IN THIS SECTION

Syntax

```
show virtual-chassis protocol database
<brief | detail | extensive>
<all-members | local | member member-id>
```

Description

Display the Virtual Chassis Control Protocol (VCCP) database statistics for all hardware devices within the Virtual Chassis or VCF.

Options

none	Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.
brief | detail | extensive	(Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` option provides more output than the `brief` option. The `extensive` option provides all output and is most useful for customer support personnel.
all-members	(Optional) Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.
local	(Optional) Display VCCP database statistics for the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display VCCP database statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 15 on page 327 lists the output fields for the `show virtual-chassis protocol database` command. Output fields are listed in the approximate order in which they appear.

Table 15: show virtual-chassis protocol database Output Fields

Field Name	Field Description	Level of Output
`LSP ID`	Link-state protocol (LSP) data unit identifier.	All levels
`Sequence`	Sequence number of the LSP.	All levels
`Checksum`	Checksum value of the LSP.	All levels
`Lifetime`	Remaining lifetime of the LSP, in seconds.	All levels
`Neighbor`	MAC address of the neighbor on the advertising system.	detail
`Interface`	Virtual Chassis port (VCP) interface name.	detail
`Metric`	Metric of the prefix or neighbor.	detail

The `extensive` output was omitted from this list. The `extensive` output is useful for customer support personnel only.

Sample Output

show virtual-chassis protocol database

```
user@switch> show virtual-chassis protocol database
member0:
--------------------------------------------------------------------------
LSP ID                        Sequence Checksum Lifetime
0000.4a75.9b7c.00-00           0x1dd80   0xc2e3      116
```

```
0000.73e9.9a57.00-00          0xf361   0x27e8      113
0021.59f7.d000.00-00         0x16882   0x3993      118
0026.888d.6800.00-00         0x1691f   0x82b7      116
  4 LSPs

member1:
--------------------------------------------------------------------------
LSP ID                       Sequence Checksum Lifetime
0000.4a75.9b7c.00-00         0x1dd80   0xc2e3      116
0000.73e9.9a57.00-00          0xf361   0x27e8      114
0021.59f7.d000.00-00         0x16883    0x289      116
0026.888d.6800.00-00         0x1691f   0x82b7      118
  4 LSPs

member8:
--------------------------------------------------------------------------
LSP ID                       Sequence Checksum Lifetime
0000.4a75.9b7c.00-00         0x1dd80   0xc2e3      118
0000.73e9.9a57.00-00          0xf361   0x27e8      114
0021.59f7.d000.00-00         0x16883    0x289      116
0026.888d.6800.00-00         0x16920   0xa335      116
  4 LSPs

member9:
--------------------------------------------------------------------------
LSP ID                       Sequence Checksum Lifetime
0000.4a75.9b7c.00-00         0x1dd80   0xc2e3      116
0000.73e9.9a57.00-00          0xf361   0x27e8      116
0021.59f7.d000.00-00         0x16883    0x289      114
0026.888d.6800.00-00         0x16920   0xa335      116
  4 LSPs
```

show virtual-chassis protocol database detail

```
user@switch> show virtual-chassis protocol database detail
member0:
--------------------------------------------------------------------------


0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-1/1.32768  Metric:     150
```

```
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/2.32768  Metric:     150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 114 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/1.32768  Metric:     150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 118 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-4/0/1.32768 Metric:      15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-0/1.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-3/0/4.32768 Metric:      15

member1:
--------------------------------------------------------------------------


0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-1/1.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/2.32768  Metric:     150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 116 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/1.32768  Metric:     150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-4/0/1.32768 Metric:      15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 117 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-0/1.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-3/0/4.32768 Metric:      15

member8:
--------------------------------------------------------------------------


0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 118 secs
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-1/1.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/2.32768  Metric:     150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 115 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/1.32768  Metric:     150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-4/0/1.32768 Metric:      15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-0/1.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-3/0/4.32768 Metric:      15

member9:
--------------------------------------------------------------------------


0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 116 secs
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-1/1.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/2.32768  Metric:     150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 117 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-1/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-1/1.32768  Metric:     150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 113 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0026.888d.6800.00  Interface: vcp-4/0/1.32768 Metric:      15

0026.888d.6800.00-00 Sequence: 0x1694f, Checksum: 0xa61a, Lifetime: 116 secs
   Neighbor: 0000.4a75.9b7c.00  Interface: vcp-0/0.32768  Metric:     150
   Neighbor: 0000.73e9.9a57.00  Interface: vcp-0/1.32768  Metric:     150
   Neighbor: 0021.59f7.d000.00  Interface: vcp-3/0/4.32768 Metric:      15
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

Understanding Virtual Chassis Components

show virtual-chassis protocol interface

IN THIS SECTION

Syntax

```
show virtual-chassis protocol interface
<brief | detail>
<all-members | local | member member-id>
<interface-name>
```

Description

Display information about Virtual Chassis Control Protocol (VCCP) statistics for VCCP-enabled interfaces within the Virtual Chassis or VCF.

Options

none	Display the VCCP interface statistics in brief form for all members of the Virtual Chassis or VCF.
brief | detail	(Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` option provides more output than the `brief` option.
all-members	(Optional) Display VCCP interface statistics for all members of the Virtual Chassis or VCF.
interface-name	(Optional) Display VCCP interface statistics for the specified interface.
local	(Optional) Display VCCP interface statistics for the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display VCCP interface statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 16 on page 332 lists the output fields for the `show virtual-chassis protocol interface` command. Output fields are listed in the approximate order in which they appear.

Table 16: show virtual-chassis protocol interface Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the VCP.	All levels

Table 16: show virtual-chassis protocol interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the link. Outputs include: • Up—The link is up. • Down—The link is down.	All levels
Metric	Metric of the prefix or neighbor.	All levels

Sample Output

show virtual-chassis protocol interface

```
user@switch> show virtual-chassis protocol interface
member0:
--------------------------------------------------------------------------
IS-IS interface database:
Interface               State         Metric
vcp-0/0.32768           Up             150
vcp-0/1.32768           Up             150
vcp-4/0/1.32768         Up             15
vcp-4/0/7.32768         Down           15

member1:
--------------------------------------------------------------------------
IS-IS interface database:
Interface               State         Metric
vcp-0/0.32768           Up             150
vcp-0/1.32768           Up             150
vcp-3/0/4.32768         Up             15

member8:
--------------------------------------------------------------------------
IS-IS interface database:
Interface               State         Metric
```

```
vcp-0/0.32768          Down           150
vcp-1/0.32768          Up             150
vcp-1/1.32768          Up             150
vcp-1/2.32768          Up             150
vcp-1/3.32768          Down           150
vcp-2/0.32768          Up             150
vcp-2/1.32768          Down           150
vcp-2/2.32768          Down           150
vcp-2/3.32768          Down           150

member9:
--------------------------------------------------------------------------
IS-IS interface database:
Interface              State         Metric
vcp-0/0.32768          Down           150
vcp-1/0.32768          Up             150
vcp-1/1.32768          Up             150
vcp-1/2.32768          Down           150
vcp-1/3.32768          Down           150
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

- Understanding EX Series Virtual Chassis
- Understanding QFX Series Virtual Chassis
- Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis
- Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

show virtual-chassis protocol route

IN THIS SECTION

Syntax

```
show virtual-chassis protocol route
<all-members | local | member member-id>
<destination-id>
```

Description

Display the unicast and multicast Virtual Chassis Control Protocol (VCCP) routing tables within the Virtual Chassis or VCF.

Options

none	Display the unicast and multicast routing tables for all members of the Virtual Chassis.

all-members	(Optional) Display the unicast and multicast routing tables for all members of the Virtual Chassis or VCF.
destination-id	(Optional) Display the unicast and multicast routing tables to the specified destination member ID for each member of the Virtual Chassis or VCF.
local	(Optional) Display the unicast and multicast routing tables on the device where you enter this command.
member *member-id*	(Optional) Display the unicast and multicast routing tables for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 17 on page 336 lists the output fields for the `show virtual-chassis protocol route` command. Output fields are listed in the approximate order in which they appear.

Table 17: show virtual-chassis protocol route Output Fields

Field Name	Field Description
`Dev`	MAC address of the member storing the VCCP routing table.
`Version`	Version of the shortest-path-first algorithm that generated the routing table.
`System ID`	MAC address of the device.
`Version`	Version of the shortest-path-first (SPF) algorithm that generated the route.
`Metric`	The metric number to get to that device.

Table 17: show virtual-chassis protocol route Output Fields *(Continued)*

Field Name	Field Description
`Interface`	Name of the Virtual Chassis port (VCP) interface connecting the devices.
`Via`	MAC address of the next-hop device, if applicable.

Sample Output

show virtual-chassis protocol route

```
user@switch> show virtual-chassis protocol route
member0:
--------------------------------------------------------------------------

Dev 0021.59f7.d000 ucast routing table             Current version: 21
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          21      150 vcp-0/1.32768 0000.4a75.9b7c
0000.73e9.9a57          21      165 vcp-4/0/1.32768 0026.888d.6800
0021.59f7.d000          21        0
0026.888d.6800          21       15 vcp-4/0/1.32768 0026.888d.6800

Dev 0021.59f7.d000 mcast routing table             Current version: 21
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          21
0000.73e9.9a57          21
0021.59f7.d000          21          vcp-4/0/1.32768
                                    vcp-0/1.32768
0026.888d.6800          21

member1:
--------------------------------------------------------------------------

Dev 0026.888d.6800 ucast routing table             Current version: 25
```

```
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          25      150 vcp-0/0.32768 0000.4a75.9b7c
0000.73e9.9a57          25      150 vcp-0/1.32768 0000.73e9.9a57
0021.59f7.d000          25       15 vcp-3/0/4.32768 0021.59f7.d000
0026.888d.6800          25        0

Dev 0026.888d.6800 mcast routing table             Current version: 25
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          25
0000.73e9.9a57          25          vcp-3/0/4.32768
0021.59f7.d000          25          vcp-0/1.32768
0026.888d.6800          25          vcp-3/0/4.32768
                                    vcp-0/0.32768
                                    vcp-0/1.32768

member8:
--------------------------------------------------------------------------

Dev 0000.4a75.9b7c ucast routing table             Current version: 39
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          39        0
0000.73e9.9a57          39      150 vcp-1/0.32768 0000.73e9.9a57
0021.59f7.d000          39      150 vcp-2/0.32768 0021.59f7.d000
0026.888d.6800          39      150 vcp-1/2.32768 0026.888d.6800

Dev 0000.4a75.9b7c mcast routing table             Current version: 39
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          39          vcp-1/0.32768
                                    vcp-2/0.32768
                                    vcp-1/2.32768
0000.73e9.9a57          39
0021.59f7.d000          39
0026.888d.6800          39

member9:
--------------------------------------------------------------------------

Dev 0000.73e9.9a57 ucast routing table             Current version: 31
----------------
```

```
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          31      150 vcp-1/0.32768 0000.4a75.9b7c
0000.73e9.9a57          31        0
0021.59f7.d000          31      165 vcp-1/1.32768 0026.888d.6800
0026.888d.6800          31      150 vcp-1/1.32768 0026.888d.6800

Dev 0000.73e9.9a57 mcast routing table             Current version: 31
----------------
System ID          Version   Metric Interface     Via
0000.4a75.9b7c          31
0000.73e9.9a57          31          vcp-1/0.32768
                                    vcp-1/1.32768
0021.59f7.d000          31
0026.888d.6800          31
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding EX Series Virtual Chassis

Understanding QFX Series Virtual Chassis

Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

show virtual-chassis protocol statistics

IN THIS SECTION

Syntax

```
show virtual-chassis protocol statistics
<all-members | local | member member-id>
<interface-name>
```

Description

Display the Virtual Chassis Control Protocol (VCCP) statistics for all hardware devices within the Virtual Chassis or VCF.

Options

none	Display VCCP statistics for all members of the Virtual Chassis or VCF.
all-members	(Optional) Display VCCP statistics for all members of the Virtual Chassis or VCF.
interface-name	(Optional) Display VCCP statistics for the specified interface.
local	(Optional) Display VCCP statistics for the switch or external Routing Engine where you run this command.
member *member-id*	(Optional) Display VCCP statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Output Fields

Table 18 on page 341 lists the output fields for the `show virtual-chassis protocol interface` command. Output fields are listed in the approximate order in which they appear.

Table 18: show virtual-chassis protocol statistics Output Fields

Field Name	Field Description
`PDU type`	Protocol data unit type.
`Received`	Number of PDUs received since VCCP started or since the statistics were set to zero.
`Processed`	Number of PDUs received minus the number of PDUs dropped.
`Drops`	Number of PDUs dropped.
`Sent`	Number of PDUs transmitted since VCCP started or since the statistics were set to zero.
`Rexmit`	Number of PDUs retransmitted since VCCP started or since the statistics were set to zero.
`Total Packets Received`	Number of PDUs received since VCCP started or since the statistics were set to zero.
`Total Packets Sent`	Number of PDUs sent since VCCP started or since the statistics were set to zero.
`LSP queue length`	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
`SPF runs`	Number of shortest-path-first (SPF) calculations that have been performed.

Table 18: show virtual-chassis protocol statistics Output Fields *(Continued)*

Field Name	Field Description
Fragments Rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP Regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software determines that a link-state PDU must be removed from the network.

Sample Output

show virtual-chassis protocol statistics

```
user@switch> show virtual-chassis protocol statistics
member0:
--------------------------------------------------------------------------


IS-IS statistics for 0021.59f7.d000:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP                8166         8166            0         4551            0
HELLO              1659         1659            0         1693            0
CSNP                  2            2            0            3            0
PSNP               1909         1909            0         2293            0
Unknown               0            0            0            0            0
Totals            11736        11736            0         8540            0

Total packets received: 11736 Sent: 8540

LSP queue length: 0 Drops: 0
SPF runs: 9
Fragments rebuilt: 1640
LSP regenerations: 1
Purges initiated: 0
```

```
member1:
--------------------------------------------------------------------------

IS-IS statistics for 0026.888d.6800:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP               10909        10909            0        12088            0
HELLO              1877         1877            0         2251            0
CSNP                  3            3            0            3            0
PSNP               3846         3846            0         3732            0
Unknown               0            0            0            0            0
Totals            16635        16635            0        18074            0

Total packets received: 16635 Sent: 18074

LSP queue length: 0 Drops: 0
SPF runs: 13
Fragments rebuilt: 1871
LSP regenerations: 2
Purges initiated: 0

member8:
--------------------------------------------------------------------------

IS-IS statistics for 0000.4a75.9b7c:
PDU type       Received    Processed        Drops         Sent       Rexmit
LSP                7935         7935            0        14865            0
HELLO              2695         2695            0         7124            0
CSNP                  4            4            0            4            0
PSNP               4398         4398            0         3666            0
Unknown               0            0            0            0            0
Totals            15032        15032            0        25659            0

Total packets received: 15032 Sent: 25659

LSP queue length: 0 Drops: 0
SPF runs: 26
Fragments rebuilt: 2666
LSP regenerations: 4
Purges initiated: 0

member9:
--------------------------------------------------------------------------
```

```
IS-IS statistics for 0000.73e9.9a57:
PDU type       Received    Processed       Drops        Sent      Rexmit
LSP               10800        10800           0        6327           0
HELLO              1492         1492           0        2356           0
CSNP                  2            2           0           2           0
PSNP               2683         2683           0        3149           0
Unknown               0            0           0           0           0
Totals            14977        14977           0       11834           0

Total packets received: 14977 Sent: 11834

LSP queue length: 0 Drops: 0
SPF runs: 19
Fragments rebuilt: 1510
LSP regenerations: 6
Purges initiated: 0
```

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding EX Series Virtual Chassis

Understanding QFX Series Virtual Chassis

Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis

show virtual-chassis vc-path

IN THIS SECTION

Syntax

```
show virtual-chassis vc-path source-interface interface-name destination-interface interface-name
```

Description

Show the forwarding path a packet takes when going from a source interface to a destination interface in a Virtual Chassis or VCF configuration.

Starting in Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, when used in VCF configurations, this command displays additional information about next hops, including when the forwarding path has multiple possible next hops within the VCF.

Options

source-interface *interface-name*	Name of the interface from which the packet originates in the Virtual Chassis or VCF
destination-interface *interface-name*	Name of the interface to which the packet is being delivered in the Virtual Chassis or VCF

Required Privilege Level

view

Output Fields

Table 19 on page 346 lists the output fields for the `show virtual-chassis vc-path` command. Output fields are listed in the approximate order in which they appear.

Some output field names and display order differ between the output for a Virtual Chassis and for a VCF. When this command is used to display forwarding paths in a VCF, additional fields are included in the output to show details of multiple possible next hops. The differences are described in Table 19 on page 346 and shown in sample output for each mode.

Table 19: show virtual-chassis vc-path Output Fields

Field Name	Field Description
`Hop`	The hop number along the path between the source and destination interfaces. The first hop entry (`Hop 0`) is the packet's source, intermediate hop information represents transitions through the members within the Virtual Chassis or VCF, and the last hop entry represents arrival at the packet's destination. For VCF output, each `Hop` entry also shows information about multiple possible next hops towards the destination. See `Next-hop PFE`, `Interface`, and `Bandwidth` output field descriptions for details.
`Member` (Virtual Chassis) `Member-ID` (VCF)	The Virtual Chassis or VCF member ID of the switch that contains the Packet Forwarding Engine for each hop through which the packet passes.
`PFE-Device`(Virtual Chassis) `PFE`(VCF)	The number of the Packet Forwarding Engine in each Virtual Chassis or VCF member through which a packet passes. For Virtual Chassis output, the Packet Forwarding Engine in each row is the next hop of the preceding Packet Forwarding Engine, including intermediate transitions through members within the Virtual Chassis. VCF output is expanded on multiple rows to show more information about multiple possible next hops—each hop entry is followed by one or more `Next-hop PFE` output fields with interface name and bandwidth information for each possible next hop.

Table 19: show virtual-chassis vc-path Output Fields *(Continued)*

Field Name	Field Description
`Next-hop PFE` (VCF only)	One or more possible next-hop Packet Forwarding Engine numbers for VCF member `Member-ID`. Each `Next-hop` entry includes the next-hops's interface name and bandwidth.
`Interface`	The name of the interface through which the Packet Forwarding Engines are connected. The interface for the first hop (`Hop 0`) is always the source interface. VCF output provides more information about multiple possible next hops for each hop entry, listing the `Interface` and the interface `Bandwidth` for each possible `Next-hop PFE`.
`Bandwidth` (VCF only)	The bandwidth (in Gbps) of the next-hop interface for the associated `Next-hop PFE` entry.

Sample Output

show virtual-chassis vc-path source-interface destination-interface (Virtual Chassis)

```
user@switch> show virtual-chassis vc-path source-interface ge-0/0/0 destination-interface
ge-1/0/1
vc-path from ge-0/0/0 to ge-1/0/1
Hop      Member     PFE-Device      Interface
0         0          1              ge-0/0/0
1         0          0              internal-1/24
2         1          3              vcp-0
3         1          4              ge-1/0/1
```

show virtual-chassis vc-path source-interface destination-interface (Virtual Chassis Fabric)

This example shows `vc-path` command output for two source-to-destination paths in the following Virtual Chassis Fabric displayed by the `show virtual-chassis` command:

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: ec56.0915.8595
Fabric Mode: Enabled
                                             Mstr         Mixed Route Neighbor List
Member ID Status Serial No    Model          prio  Role    Mode  Mode  ID  Interface
1 (FPC 1) Prsnt  ...          qfx5100-24q-aa 129   Backup     N   F     7  vcp-255/0/7
                                                                        8  vcp-255/0/8
                                                                        12 vcp-255/0/12
                                                                        3  vcp-255/0/3
                                                                        4  vcp-255/0/4
                                                                        5  vcp-255/0/5
                                                                        6  vcp-255/0/6
                                                                        10 vcp-255/0/10
                                                                        11 vcp-255/0/11
                                                                        9  vcp-255/0/9
2 (FPC 2) Prsnt  ...          qfx5100-24q-aa 129   Master*    N   F     7  vcp-255/0/7
                                                                        8  vcp-255/0/8
                                                                        12 vcp-255/0/12
                                                                        3  vcp-255/0/3
                                                                        4  vcp-255/0/4
                                                                        5  vcp-255/0/5
                                                                        6  vcp-255/0/6
                                                                        10 vcp-255/0/10
                                                                        11 vcp-255/0/11
                                                                        9  vcp-255/0/9
3 (FPC 3)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N   F    1  vcp-255/0/1
                                                                        2  vcp-255/0/2
4 (FPC 4)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N   F    1  vcp-255/0/1
                                                                        2  vcp-255/0/2
5 (FPC 5)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N   F    1  vcp-255/0/1
                                                                        2  vcp-255/0/2
6 (FPC 6)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N   F    1  vcp-255/0/1
                                                                        2  vcp-255/0/2
7 (FPC 7)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N   F    1  vcp-255/0/1
                                                                        2  vcp-255/0/2
```

```
8 (FPC 8)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
9 (FPC 9)  Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
10 (FPC 10)Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
11 (FPC 11)Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2
12 (FPC 12)Prsnt  ...          qfx5100-24q-aa   0   Linecard   N  F   1  vcp-255/0/1
                                                                      2  vcp-255/0/2


user@switch> show virtual-chassis vc-path source-interface xe-12/0/25 destination-interface
xe-5/0/25
Fabric forwarding path from xe-12/0/25 (PFE 12) to xe-5/0/25 (PFE 5)
  Hop   0  Member-ID  12  PFE   12
    Next-hop PFE   1
      Interface vcp-255/0/1.32768  Bandwidth   40
    Next-hop PFE   2
      Interface vcp-255/0/2.32768  Bandwidth   40
  Hop   1  Member-ID   1  PFE    1
    Next-hop PFE   5
      Interface vcp-255/0/5.32768  Bandwidth   40
  Hop   1  Member-ID   2  PFE    2
    Next-hop PFE   5
      Interface vcp-255/0/5.32768  Bandwidth   40
  Hop   2  Member-ID   5  PFE    5


user@switch> show virtual-chassis vc-path source-interface xe-12/0/25 destination-interface
xe-1/0/25
Fabric forwarding path from xe-12/0/25 (PFE 12) to xe-1/0/25 (PFE 1)
  Hop   0  Member-ID  12  PFE   12
    Next-hop PFE   1
      Interface vcp-255/0/1.32768  Bandwidth   40
  Hop   1  Member-ID   1  PFE    1
```

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis
Understanding EX Series Virtual Chassis
Understanding QFX Series Virtual Chassis
EX8200 Virtual Chassis Overview

show virtual-chassis vc-port

IN THIS SECTION

Syntax

```
show virtual-chassis vc-port
<all-members | local | member member-id>
```

Description

Display the status of the Virtual Chassis ports (VCPs), including dedicated VCPs, default-configured VCPs, and uplink ports configured as VCPs, if present.

Options

none	Display the operational status of all VCPs of the member switch where you enter the command.
all-members	(Optional) Display the operational status of all VCPs on all members of the Virtual Chassis or VCF.
local	(Optional) Display the operational status of the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display the operational status of all VCPs for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

view

Output Fields

Table 20 on page 351 lists the output fields for the `show virtual-chassis vc-port` command. Output fields are listed in the approximate order in which they appear.

Table 20: show virtual-chassis vc-port Output Fields

Field Name	Field Description
member name	Device or member number in the Virtual Chassis, represented as `fpc`*number*, `member`*number*, or local-re (the primary member).

Table 20: show virtual-chassis vc-port Output Fields *(Continued)*

Field Name	Field Description
`Interface or PIC/ Port`	VCP name. • The dedicated VCPs in an EX4200 or EX4500 Virtual Chassis are `vcp-0` and `vcp-1`. The dedicated VCPs in an EX4550 Virtual Chassis are `VCP-1/0`, `VCP-1/1`, `VCP-2/0`, and `VCP-2/1`. • Optical ports set as VCPs are named `1/0` and `1/1`, representing the PIC number and the port number. • The native VCP (port 0) on an XRE200 External Routing Engine in an EX8200 Virtual Chassis is named `vcp-0`. • The VCPs on each Virtual Chassis Control Interface (VCCI) module in an XRE200 External Routing Engine are named using the `vcp-`*slot-number*/*port-number* convention; for instance, `vcp-1/0`. • The VCPs on EX8200 member switches are named using the `vcp-`*slot-number*/*pic-number*/*interface-number* convention; for instance, `vcp-3/0/2`. • A `255` as the first number in your port number indicates that your VCP is part of a Link Aggregation group (LAG) bundle. For example, `vcp-255/1/0` indicates that the dedicated VCP named `vcp-1/0` is part of a LAG bundle, and `vcp-255/1/0` represents an uplink port previously named `xe-0/1/0` that is now part of a VCP LAG bundle.
`Type`	Type of VCP: • `Dedicated`—Rear panel dedicated VCP on an EX4200 or EX4300 multigigabit model switch, a Virtual Chassis module port on EX4500 or EX4550 switch, or any VCP link connected to an XRE200 External Routing Engine in an EX8200 Virtual Chassis. • `Configured`—Optical port configured as a VCP, which includes default-configured VCPs (set in the default factory configuration) or those that are set by a user. • `Auto-Configured`—Optical port automatically converted into a VCP. See Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port or Configuring an EX4650 or a QFX Series Virtual Chassis for information about configuring VCPs, and Automatic Virtual Chassis Port (VCP) Conversion for details on how a port is automatically converted into a VCP.

Table 20: show virtual-chassis vc-port Output Fields *(Continued)*

Field Name	Field Description
`Trunk ID`	A positive-number ID assigned to a link aggregation group (LAG) formed by the Virtual Chassis. The trunk ID value is -1 if no trunk is formed. A LAG between uplink VCPs requires that the link speed be the same on connected interfaces and that at least two VCPs on one member be connected to at least two VCPs on the other member in an EX4200 or EX4500 Virtual Chassis. Dedicated VCP LAGs are assigned trunk IDs 1 and 2. Trunk IDs for LAGs formed with uplink VCPs therefore have values of 3 or greater. The trunk ID value changes if the link-adjacency state between LAG members changes; trunk membership is then allocated or deallocated.
`Status`	Interface status: • `absent`—Interface is not a VCP link. • `down`—VCP link is down. • `up`—VCP link is up.
`Speed (mbps)`	Speed of the interface in megabits per second.
`Neighbor ID/ Interface`	The Virtual Chassis member ID and interface of a VCP on a member that is connected to the interface or PIC/Port field in the same row as this interface.

Sample Output

show virtual-chassis vc-port (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis vc-port

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status     Speed        Neighbor
or                             ID               (mbps)       ID  Interface
```

```
PIC / Port
vcp-0       Dedicated        1    Up        32000       1    vcp-1
vcp-1       Dedicated        2    Up        32000       0    vcp-0
1/0         Auto-Configured  3    Up        1000        2    vcp-255/1/0
1/0         Auto-Configured  3    Up        1000        2    vcp-255/1/1
```

show virtual-chassis vc-port (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis vc-port
member0:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status    Speed        Neighbor
or                                 ID              (mbps)       ID  Interface
Slot/PIC/Port
vcp-0/0         Dedicated          -1    Up        1000         8   vcp-1/1
vcp-0/1         Dedicated          -1    Up        1000         8   vcp-2/0
4/0/4           Configured         -1    Up        10000        1   vcp-3/0/4
4/0/7           Configured         -1    Down      10000
4/0/3           Configured               Absent
4/0/2           Configured               Absent
4/0/5           Configured               Absent
4/0/6           Configured               Absent
4/0/1           Configured               Absent
4/0/0           Configured               Absent

member1:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status    Speed        Neighbor
or                                 ID                 (mbps)    ID  Interface
Slot/PIC/Port
vcp-0/0         Dedicated          -1    Up        1000         8   vcp-1/2
3/0/0           Configured         -1    Down      10000
3/0/1           Configured         -1    Down      10000
3/0/4           Configured         -1    Up        10000        0   vcp-4/0/4
3/0/5           Configured               Absent
4/0/5           Configured               Absent
4/0/4           Configured               Absent

member8:
--------------------------------------------------------------------------
Interface       Type              Trunk  Status    Speed        Neighbor
```

```
or                                ID            (mbps)      ID  Interface
Slot/PIC/Port
vcp-0/0         Dedicated         -1    Down      1000
vcp-1/0         Dedicated         -1    Up        1000        9   vcp-1/0
vcp-1/1         Dedicated         -1    Up        1000        0   vcp-0/0
vcp-1/2         Dedicated         -1    Up        1000        1   vcp-0/0
vcp-1/3         Dedicated         -1    Up        1000        9   vcp-1/3
vcp-2/0         Dedicated         -1    Up        1000        0   vcp-0/1
vcp-2/1         Dedicated         -1    Up        1000        9   vcp-1/2
vcp-2/2         Dedicated         -1    Down      1000
vcp-2/3         Dedicated         -1    Down         1000

member9:
--------------------------------------------------------------------------
Interface       Type             Trunk  Status    Speed       Neighbor
or                                ID            (mbps)      ID  Interface
Slot/PIC/Port
vcp-0/0         Dedicated         -1    Disabled  1000
vcp-1/0         Dedicated         -1    Up        1000        8   vcp-1/0
vcp-1/1         Dedicated         -1    Down      1000
vcp-1/2         Dedicated         -1    Up        1000        8   vcp-2/1
vcp-1/3         Dedicated         -1    Up        1000        8   vcp-1/3
```

show virtual-chassis vc-port all-members

```
user@switch> show virtual-chassis vc-port all-members



fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status  Speed    Neighbor
or                             ID            (mbps)   ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up      32000     1   vcp-1
vcp-1       Dedicated           2    Up      32000     0   vcp-0
1/0         Auto-Configured     3    Up      1000      2   vcp-255/1/0
1/1         Auto-Configured     3    Up      1000      2   vcp-255/1/1

fpc1:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status  Speed    Neighbor
```

```
or                              ID            (mbps)   ID  Interface
PIC / Port
vcp-0       Dedicated           1     Up      32000     0   vcp-1
vcp-1       Dedicated           2     Up      32000     0   vcp-0
1/0         Auto-Configured     -1    Up      1000      3   vcp-255/1/0

fpc2:
--------------------------------------------------------------------------
Interface   Type               Trunk  Status   Speed    Neighbor
or                              ID              (mbps)   ID  Interface
PIC / Port
vcp-0       Dedicated           1     Up      32000     3   vcp-1
vcp-1       Dedicated           2     Up      32000     3   vcp-0
1/0         Auto-Configured     3     Up      1000      0   vcp-255/1/0
1/1         Auto-Configured     3     Up      1000      0   vcp-255/1/1

fpc3:
--------------------------------------------------------------------------
Interface   Type               Trunk  Status   Speed    Neighbor
or                              ID              (mbps)   ID  Interface
PIC / Port
vcp-0       Dedicated           1     Up      32000     2   vcp-0
vcp-1       Dedicated           2     Up      32000     2   vcp-1
1/0         Auto-Configured     -1    Up      1000      1   vcp-255/1/0
```

show virtual-chassis vc-port (QFX5120-32C Virtual Chassis)

```
user@switch> show virtual-chassis vc-port

localre:
-
Interface   Type               Trunk  Status  Speed     Neighbor
or                              ID            (mbps)    ID  Interface
PIC / Port
0/13        Auto-Configured     6     Up      100000    1   vcp-255/0/13
0/23        Auto-Configured     6     Up      100000    1   vcp-255/0/23
0/26        Auto-Configured     5     Up      40000     1   vcp-255/0/26
0/27        Configured          5     Up      40000     1   vcp-255/0/27
0/6         Configured          5     Up      40000     1   vcp-255/0/6
0/8         Auto-Configured     5     Up      40000     1   vcp-255/0/8
0/9         Auto-Configured     6     Up      100000    1   vcp-255/0/9
```

```
fpc1:
-
Interface    Type              Trunk  Status  Speed     Neighbor
or                              ID             (mbps)    ID  Interface
PIC / Port
0/13         Auto-Configured     6    Up      100000    0   vcp-255/0/13
0/23         Auto-Configured     6    Up      100000    0   vcp-255/0/23
0/26         Auto-Configured     5    Up      40000     0   vcp-255/0/26
0/27         Configured          5    Up      40000     0   vcp-255/0/27
0/6          Configured          5    Up      40000     0   vcp-255/0/6
0/8          Auto-Configured     5    Up      40000     0   vcp-255/0/8
0/9          Auto-Configured     6    Up      100000    0   vcp-255/0/9
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

show virtual-chassis vc-port statistics

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

show virtual-chassis vc-port diagnostics optics

IN THIS SECTION

Syntax

```
show virtual-chassis vc-port diagnostics optics
<all-members | local | member member-id>
<interface-name>
```

Description

Display diagnostics data and alarms for Ethernet optical transceivers installed in ports configured as Virtual Chassis Ports (VCPs) in an EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that a transceiver is not operating properly. DOM information can be used to diagnose why a transceiver is not working.

On some EX Series switches, the `request virtual-chassis vc-port diagnostics optics` command must be entered to run a diagnostic scan before you can gather the `show virtual-chassis vc-port diagnostics optics` output.

Options

none	Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.
all-members	(Optional) Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.
interface-name	(Optional) Display diagnostics information for the transceiver installed in a specified VCP.

local	(Optional) Display diagnostics information for transceivers installed in VCPs on the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display diagnostics information for transceivers installed in VCPs on a specified member of a Virtual Chassis or VCF.

Required Privilege Level

view

Output Fields

Table 21 on page 359 lists the output fields for the `show virtual-chassis vc-port diagnostics optics` command. Output fields are listed in the approximate order in which they appear.

Table 21: show virtual-chassis vc-port diagnostics optics Output Fields

Field Name	Field Description
FPC	Displays the FPC slot number.
Virtual chassis port	Displays the name of the VCP.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.

Table 21: show virtual-chassis vc-port diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is *On* or *Off*.
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is *On* or *Off*.
Laser bias current high warning	Displays whether the laser bias power setting high warning is *On* or *Off*.
Laser bias current low warning	Displays whether the laser bias power setting low warning is *On* or *Off*.
Laser output power high alarm	Displays whether the laser output power high alarm is *On* or *Off*.
Laser output power low alarm	Displays whether the laser output power low alarm is *On* or *Off*.
Laser output power high warning	Displays whether the laser output power high warning is *On* or *Off*.
Laser output power low warning	Displays whether the laser output power low warning is *On* or *Off*.
Module temperature high alarm	Displays whether the module temperature high alarm is *On* or *Off*.
Module temperature low alarm	Displays whether the module temperature low alarm is *On* or *Off*.
Module temperature high warning	Displays whether the module temperature high warning is *On* or *Off*.
Module temperature low warning	Displays whether the module temperature low warning is *On* or *Off*.
Module voltage high alarm	Displays whether the module voltage high alarm is *On* or *Off*.
Module voltage low alarm	Displays whether the module voltage low alarm is *On* or *Off*.

Table 21: show virtual-chassis vc-port diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Module voltage high warning	Displays whether the module voltage high warning is *On* or *Off*.
Module voltage low warning	Displays whether the module voltage low warning is *On* or *Off*.
Laser rx power high alarm	Displays whether the receive laser power high alarm is *On* or *Off*.
Laser rx power low alarm	Displays whether the receive laser power low alarm is *On* or *Off*.
Laser rx power high warning	Displays whether the receive laser power high warning is *On* or *Off*.
Laser rx power low warning	Displays whether the receive laser power low warning is *On* or *Off*.
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.

Table 21: show virtual-chassis vc-port diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.

Table 21: show virtual-chassis vc-port diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

show virtual-chassis vc-port diagnostics optics

```
user@switch> show virtual-chassis vc-port diagnostics optics
fpc0:
--------------------------------------------------------------------------
Virtual chassis port: vcp-0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                       :  N/A

fpc1:
--------------------------------------------------------------------------
Virtual chassis port: vcp-0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                       :  N/A

fpc2:
--------------------------------------------------------------------------
Virtual chassis port: vcp-2/0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/15
```

```
    Optical diagnostics                        :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                         :  4.130 mA
    Laser output power                         :  0.2450 mW / -6.11 dBm
    Module temperature                         :  32 degrees C / 90 degrees F
    Module voltage                             :  3.3530 V
    Receiver signal average optical power      :  0.0971 mW / -10.13 dBm
    Laser bias current high alarm              :  Off
    Laser bias current low alarm               :  Off
    Laser bias current high warning            :  Off
    Laser bias current low warning             :  Off
    Laser output power high alarm              :  Off
    Laser output power low alarm               :  Off
    Laser output power high warning            :  Off
    Laser output power low warning             :  Off
    Module temperature high alarm              :  Off
    Module temperature low alarm               :  Off
    Module temperature high warning            :  Off
    Module temperature low warning             :  Off
    Module voltage high alarm                  :  Off
    Module voltage low alarm                   :  Off
    Module voltage high warning                :  Off
    Module voltage low warning                 :  Off
    Laser rx power high alarm                  :  Off
    Laser rx power low alarm                   :  Off
    Laser rx power high warning                :  Off
    Laser rx power low warning                 :  Off
    Laser bias current high alarm threshold    :  14.998 mA
    Laser bias current low alarm threshold     :  0.998 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold   :  1.198 mA
    Laser output power high alarm threshold    :  0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold     :  0.0790 mW / -11.02 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold   :  0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold    :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold     :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  80 degrees C / 176 degrees F
    Module temperature low warning threshold   :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold        :  3.600 V
    Module voltage low alarm threshold         :  3.000 V
    Module voltage high warning threshold      :  3.499 V
    Module voltage low warning threshold       :  3.099 V
```

```
    Laser rx power high alarm threshold       :  1.5848 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  1.2589 mW / 1.00 dBm
    Laser rx power low warning threshold      :  0.0125 mW / -19.03 dBm
  Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.428 mA
    Laser output power                        :  0.4760 mW / -3.22 dBm
    Module temperature                        :  28 degrees C / 83 degrees F
    Module voltage                            :  3.3440 V
    Receiver signal average optical power     :  0.4002 mW / -3.98 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
```

```
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm

fpc3:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/2
    Laser bias current                        :  7.876 mA
    Laser output power                        :  0.5330 mW / -2.73 dBm
    Module temperature                        :  26 degrees C / 78 degrees F
    Module voltage                            :  3.3060 V
    Receiver signal average optical power     :  0.4885 mW / -3.11 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  14.500 mA
    Laser bias current low alarm threshold    :  3.500 mA
    Laser bias current high warning threshold :  14.500 mA
    Laser bias current low warning threshold  :  3.500 mA
    Laser output power high alarm threshold   :  1.8620 mW / 2.70 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7410 mW / -1.30 dBm
```

```
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.9952 mW / 3.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.052 mA
    Laser output power                        :  0.5030 mW / -2.98 dBm
    Module temperature                        :  24 degrees C / 75 degrees F
    Module voltage                            :  3.2890 V
    Receiver signal average optical power     :  0.5028 mW / -2.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
```

```
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/4
    Laser bias current                        :  7.978 mA
    Laser output power                        :  0.5460 mW / -2.63 dBm
    Module temperature                        :  24 degrees C / 76 degrees F
    Module voltage                            :  3.3060 V
    Receiver signal average optical power     :  0.6305 mW / -2.00 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  14.500 mA
```

```
    Laser bias current low alarm threshold    :  3.500 mA
    Laser bias current high warning threshold :  14.500 mA
    Laser bias current low warning threshold  :  3.500 mA
    Laser output power high alarm threshold   :  1.8620 mW / 2.70 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7410 mW / -1.30 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.9952 mW / 3.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
fpc4:
--------------------------------------------------------------------------
Virtual chassis port: vcp-0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/4
    Laser bias current                        :  7.860 mA
    Laser output power                        :  0.5370 mW / -2.70 dBm
    Module temperature                        :  24 degrees C / 75 degrees F
    Module voltage                            :  3.2920 V
    Receiver signal average optical power     :  0.6271 mW / -2.03 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
```

```
      Module voltage high alarm                : Off
      Module voltage low alarm                 : Off
      Module voltage high warning              : Off
      Module voltage low warning               : Off
      Laser rx power high alarm                : Off
      Laser rx power low alarm                 : Off
      Laser rx power high warning              : Off
      Laser rx power low warning               : Off
      Laser bias current high alarm threshold  : 14.500 mA
      Laser bias current low alarm threshold   : 3.500 mA
      Laser bias current high warning threshold : 14.500 mA
      Laser bias current low warning threshold  : 3.500 mA
      Laser output power high alarm threshold  : 1.8620 mW / 2.70 dBm
      Laser output power low alarm threshold   : 0.0740 mW / -11.31 dBm
      Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
      Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
      Module temperature high alarm threshold  : 75 degrees C / 167 degrees F
      Module temperature low alarm threshold   : -5 degrees C / 23 degrees F
      Module temperature high warning threshold : 70 degrees C / 158 degrees F
      Module temperature low warning threshold  : 0 degrees C / 32 degrees F
      Module voltage high alarm threshold      : 3.630 V
      Module voltage low alarm threshold       : 2.970 V
      Module voltage high warning threshold    : 3.465 V
      Module voltage low warning threshold     : 3.135 V
      Laser rx power high alarm threshold      : 1.9952 mW / 3.00 dBm
      Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
      Laser rx power high warning threshold    : 0.7943 mW / -1.00 dBm
      Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm
```

show virtual-chassis vc-port diagnostics optics (interface-name)

```
user@external-routing-engine> show virtual-chassis vc-port diagnostics optics vcp-255/0/3
fpc0:
--------------------------------------------------------------------------


fpc1:
--------------------------------------------------------------------------


fpc2:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/3
```

```
    Laser bias current                         :  5.448 mA
    Laser output power                         :  0.4770 mW / -3.21 dBm
    Module temperature                         :  28 degrees C / 82 degrees F
    Module voltage                             :  3.3450 V
    Receiver signal average optical power      :  0.3973 mW / -4.01 dBm
    Laser bias current high alarm              :  Off
    Laser bias current low alarm               :  Off
    Laser bias current high warning            :  Off
    Laser bias current low warning             :  Off
    Laser output power high alarm              :  Off
    Laser output power low alarm               :  Off
    Laser output power high warning            :  Off
    Laser output power low warning             :  Off
    Module temperature high alarm              :  Off
    Module temperature low alarm               :  Off
    Module temperature high warning            :  Off
    Module temperature low warning             :  Off
    Module voltage high alarm                  :  Off
    Module voltage low alarm                   :  Off
    Module voltage high warning                :  Off
    Module voltage low warning                 :  Off
    Laser rx power high alarm                  :  Off
    Laser rx power low alarm                   :  Off
    Laser rx power high warning                :  Off
    Laser rx power low warning                 :  Off
    Laser bias current high alarm threshold    :  10.500 mA
    Laser bias current low alarm threshold     :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold   :  2.500 mA
    Laser output power high alarm threshold    :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold     :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold   :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold    :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold     :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold   :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold        :  3.630 V
    Module voltage low alarm threshold         :  2.970 V
    Module voltage high warning threshold      :  3.465 V
    Module voltage low warning threshold       :  3.135 V
    Laser rx power high alarm threshold        :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold         :  0.0407 mW / -13.90 dBm
```

```
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm

fpc3:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.040 mA
    Laser output power                        :  0.5020 mW / -2.99 dBm
    Module temperature                        :  24 degrees C / 74 degrees F
    Module voltage                            :  3.2870 V
    Receiver signal average optical power     :  0.5073 mW / -2.95 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
```

```
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm

fpc4:
--------------------------------------------------------------------------
```

show virtual-chassis vc-port diagnostics optics local

```
user@switch> show virtual-chassis vc-port diagnostics optics local
Virtual chassis port: vcp-2/0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                        :  4.130 mA
    Laser output power                        :  0.2450 mW / -6.11 dBm
    Module temperature                        :  32 degrees C / 90 degrees F
    Module voltage                            :  3.3530 V
    Receiver signal average optical power     :  0.0961 mW / -10.17 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
```

```
    Module voltage high alarm                : Off
    Module voltage low alarm                 : Off
    Module voltage high warning              : Off
    Module voltage low warning               : Off
    Laser rx power high alarm                : Off
    Laser rx power low alarm                 : Off
    Laser rx power high warning              : Off
    Laser rx power low warning               : Off
    Laser bias current high alarm threshold  : 14.998 mA
    Laser bias current low alarm threshold   : 0.998 mA
    Laser bias current high warning threshold : 14.000 mA
    Laser bias current low warning threshold  : 1.198 mA
    Laser output power high alarm threshold  : 0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold   : 0.0790 mW / -11.02 dBm
    Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  : 0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold  : 85 degrees C / 185 degrees F
    Module temperature low alarm threshold   : -10 degrees C / 14 degrees F
    Module temperature high warning threshold : 80 degrees C / 176 degrees F
    Module temperature low warning threshold  : -5 degrees C / 23 degrees F
    Module voltage high alarm threshold      : 3.600 V
    Module voltage low alarm threshold       : 3.000 V
    Module voltage high warning threshold    : 3.499 V
    Module voltage low warning threshold     : 3.099 V
    Laser rx power high alarm threshold      : 1.5848 mW / 2.00 dBm
    Laser rx power low alarm threshold       : 0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold    : 1.2589 mW / 1.00 dBm
    Laser rx power low warning threshold     : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                       : 5.426 mA
    Laser output power                       : 0.4760 mW / -3.22 dBm
    Module temperature                       : 28 degrees C / 83 degrees F
    Module voltage                           : 3.3450 V
    Receiver signal average optical power    : 0.3955 mW / -4.03 dBm
    Laser bias current high alarm            : Off
    Laser bias current low alarm             : Off
    Laser bias current high warning          : Off
    Laser bias current low warning           : Off
    Laser output power high alarm            : Off
    Laser output power low alarm             : Off
    Laser output power high warning          : Off
    Laser output power low warning           : Off
    Module temperature high alarm            : Off
```

```
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  10.500 mA
    Laser bias current low alarm threshold    :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

show virtual-chassis vc-port diagnostics optics (member member-id)

```
user@switch> show virtual-chassis vc-port diagnostics optics member 2
fpc2:
--------------------------------------------------------------------------
Virtual chassis port: vcp-2/0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-2/1
```

```
    Optical diagnostics                     :  N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                     :  N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                     :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                      :  4.130 mA
    Laser output power                      :  0.2450 mW / -6.11 dBm
    Module temperature                      :  31 degrees C / 88 degrees F
    Module voltage                          :  3.3530 V
    Receiver signal average optical power   :  0.0961 mW / -10.17 dBm
    Laser bias current high alarm           :  Off
    Laser bias current low alarm            :  Off
    Laser bias current high warning         :  Off
    Laser bias current low warning          :  Off
    Laser output power high alarm           :  Off
    Laser output power low alarm            :  Off
    Laser output power high warning         :  Off
    Laser output power low warning          :  Off
    Module temperature high alarm           :  Off
    Module temperature low alarm            :  Off
    Module temperature high warning         :  Off
    Module temperature low warning          :  Off
    Module voltage high alarm               :  Off
    Module voltage low alarm                :  Off
    Module voltage high warning             :  Off
    Module voltage low warning              :  Off
    Laser rx power high alarm               :  Off
    Laser rx power low alarm                :  Off
    Laser rx power high warning             :  Off
    Laser rx power low warning              :  Off
    Laser bias current high alarm threshold  :  14.998 mA
    Laser bias current low alarm threshold   :  0.998 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  1.198 mA
    Laser output power high alarm threshold  :  0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold   :  0.0790 mW / -11.02 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold  :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold   :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  80 degrees C / 176 degrees F
    Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
```

```
    Module voltage high alarm threshold      :  3.600 V
    Module voltage low alarm threshold       :  3.000 V
    Module voltage high warning threshold    :  3.499 V
    Module voltage low warning threshold     :  3.099 V
    Laser rx power high alarm threshold      :  1.5848 mW / 2.00 dBm
    Laser rx power low alarm threshold       :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold    :  1.2589 mW / 1.00 dBm
    Laser rx power low warning threshold     :  0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                       :  5.418 mA
    Laser output power                       :  0.4770 mW / -3.21 dBm
    Module temperature                       :  28 degrees C / 83 degrees F
    Module voltage                           :  3.3450 V
    Receiver signal average optical power    :  0.3964 mW / -4.02 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser output power high alarm            :  Off
    Laser output power low alarm             :  Off
    Laser output power high warning          :  Off
    Laser output power low warning           :  Off
    Module temperature high alarm            :  Off
    Module temperature low alarm             :  Off
    Module temperature high warning          :  Off
    Module temperature low warning           :  Off
    Module voltage high alarm                :  Off
    Module voltage low alarm                 :  Off
    Module voltage high warning              :  Off
    Module voltage low warning               :  Off
    Laser rx power high alarm                :  Off
    Laser rx power low alarm                 :  Off
    Laser rx power high warning              :  Off
    Laser rx power low warning               :  Off
    Laser bias current high alarm threshold  :  10.500 mA
    Laser bias current low alarm threshold   :  2.000 mA
    Laser bias current high warning threshold :  9.000 mA
    Laser bias current low warning threshold  :  2.500 mA
    Laser output power high alarm threshold  :  1.4120 mW / 1.50 dBm
    Laser output power low alarm threshold   :  0.0740 mW / -11.31 dBm
    Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
    Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
    Module temperature high alarm threshold  :  75 degrees C / 167 degrees F
```

```
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.630 V
    Module voltage low alarm threshold        :  2.970 V
    Module voltage high warning threshold     :  3.465 V
    Module voltage low warning threshold      :  3.135 V
    Laser rx power high alarm threshold       :  1.5849 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0407 mW / -13.90 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.1023 mW / -9.90 dBm
```

Release Information

Command introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

show virtual-chassis vc-port

Install a Transceiver

Remove a Transceiver

Junos OS Ethernet Interfaces Configuration Guide

show virtual-chassis vc-port statistics

IN THIS SECTION

Syntax

```
show virtual-chassis vc-port statistics
<all-members | local | member member-id>
<brief | detail | extensive >
<interface-name>
```

Description

Display the traffic statistics collected on Virtual Chassis ports (VCPs).

Options

none	Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.
brief | detail | extensive	(Optional) Display the specified level of output. Using the `brief` option is equivalent to entering the command with no options (the default). The `detail` and `extensive` options provide identical displays.
all-members	(Optional) Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF
interface-name	(Optional) Display traffic statistics for the specified VCP.
local	(Optional) Display traffic statistics for VCPs on the switch or external Routing Engine where you enter this command.
member *member-id*	(Optional) Display traffic statistics for VCPs on the specified member of a Virtual Chassis or VCF.

Required Privilege Level

view

Output Fields

Table 22 on page 380 lists the output fields for the `show virtual-chassis vc-port statistics` command. Output fields are listed in the approximate order in which they appear.

Table 22: show virtual-chassis vc-port statistics Output Fields

Field Name	Field Description	Level of Output
`fpc`*number*	(All Virtual Chassis except EX8200 Virtual Chassis. VCF) ID of the Virtual Chassis member. The FPC number is the same as the member ID.	All levels
`member` *number*	(EX8200 Virtual Chassis only) Member ID of the Virtual Chassis member.	All levels
`Interface`	VCP name.	`brief`
`Input Octets/Packets`	Number of octets and packets received on the VCP.	`brief`, `member`, none
`Output Octets/ Packets`	Number of octets and packets transmitted on the VCP.	`brief`, `member`, none
`master:` *number*	Member ID of the primary Routing Engine.	All levels
`Port`	VCP for which `RX` (Receive) statistics,`TX` (Transmit) statistics, or both are reported by the VCP subsystem during a sampling interval—since the statistics counter was last cleared.	`detail`, `extensive`
`Total octets`	Total number of octets received and transmitted on the VCP.	`detail`, `extensive`

Table 22: show virtual-chassis vc-port statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Total packets`	Total number of packets received and transmitted on the VCP.	`detail`, `extensive`
`Unicast packets`	Number of unicast packets received and transmitted on the VCP.	`detail`, `extensive`
`Broadcast packets`	Number of broadcast packets received and transmitted on the VCP.	`detail`, `extensive`
`Multicast packets`	Number of multicast packets received and transmitted on the VCP.	`detail`, `extensive`
`MAC control frames`	Number of media access control (MAC) control frames received and transmitted on the VCP.	`detail`, `extensive`
`CRC alignment errors`	Number of packets received on the VCP that had a length—excluding framing bits, but including frame check sequence (FCS) octets—of between 64 and 1518 octets, inclusive, and had one of the following errors: • Invalid FCS with an integral number of octets (FCS error) • Invalid FCS with a nonintegral number of octets (alignment error)	`detail`, `extensive`
`Oversize packets`	Number of packets received on the VCP that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed.	`detail`, `extensive`
`Undersize packets`	Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits but including FCS octets) and were otherwise well formed..	`detail`, `extensive`

Table 22: show virtual-chassis vc-port statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Jabber packets`	Number of packets received on the VCP that were longer than 1518 octets—excluding framing bits, but including FCS octets—and that had either an FCS error or an alignment error. **NOTE**: This definition of *jabber* is different from the definition in IEEE-802.3 section 8.2.1.5 (10Base5) and section 10.3.1.4 (10Base2). These documents define *jabber* as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.	`detail`, `extensive`
`Fragments received`	Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.	`detail`, `extensive`
`Ifout errors`	Number of outbound packets received on the VCP that could not be transmitted because of errors.	`detail`, `extensive`
`Packet drop events`	Number of outbound packets received on the VCP that were dropped, rather than being encapsulated and sent out of the switch as fragments. The packet drop counter is incremented if a temporary shortage of packet memory causes packet fragmentation to fail.	`detail`, `extensive`
`64 octets frames`	Number of packets received on the VCP (including invalid packets) that were 64 octets in length (excluding framing bits, but including FCS octets).	`detail`, `extensive`

Table 22: show virtual-chassis vc-port statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`65-127 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 65 and 127 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`128-255 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 128 and 255 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`256-511 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 256 and 511 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`512-1023 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 512 and 1023 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`1024-1518 octets frames`	Number of packets received on the VCP (including invalid packets) that were between 1024 and 1518 octets in length, inclusive (excluding framing bits, but including FCS octets).	`detail`, `extensive`
`Rate packets per second`	Number of packets per second received and transmitted on the VCP.	`detail`, `extensive`
`Rate bytes per second`	Number of bytes per second received and transmitted on the VCP.	`detail`, `extensive`

Sample Output

show virtual-chassis vc-port statistics

```
user@switch> show virtual-chassis vc-port statistics
fpc0:
--------------------------------------------------------------------------
Interface               Input Octets/Packets          Output Octets/Packets
internal-0/24           0           / 0               0           / 0
internal-0/25           0           / 0               0           / 0
internal-1/26           0           / 0               0           / 0
internal-1/27           0           / 0               0           / 0
vcp-0                   0           / 0               0           / 0
vcp-1                   0           / 0               0           / 0
internal-0/26           0           / 0               0           / 0
internal-0/27           0           / 0               0           / 0
internal-1/24           0           / 0               0           / 0
internal-1/25           0           / 0               0           / 0

{master:0}
```

show virtual-chassis vc-port statistics (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis vc-port statistics
member0:
--------------------------------------------------------------------------
Interface               Input Octets/Packets          Output Octets/Packets
vcp-4/0/4               43171238    / 48152           47687133    / 51891
vcp-4/0/7               0           / 0               0           / 0

member1:
--------------------------------------------------------------------------
Interface               Input Octets/Packets          Output Octets/Packets
vcp-3/0/0               0           / 0               0           / 0
vcp-3/0/1               0           / 0               0           / 0
vcp-3/0/4               47695376    / 51899           43180556    / 48160

member8:
--------------------------------------------------------------------------
```

```
member9:
--------------------------------------------------------------------------
```

show virtual-chassis vc-port statistics brief

```
user@switch> show virtual-chassis vc-port statistics brief
fpc0:
--------------------------------------------------------------------------
Interface              Input Octets/Packets          Output Octets/Packets
internal-0/24          0           / 0               0           / 0
internal-0/25          0           / 0               0           / 0
internal-1/26          0           / 0               0           / 0
internal-1/27          0           / 0               0           / 0
vcp-0                  0           / 0               0           / 0
vcp-1                  0           / 0               0           / 0
internal-0/26          0           / 0               0           / 0
internal-0/27          0           / 0               0           / 0
internal-1/24          0           / 0               0           / 0
internal-1/25          0           / 0               0           / 0

{master:0}
```

show virtual-chassis vc-port statistics extensive

```
user@switch> show virtual-chassis vc-port statistics extensive
fpc0:
--------------------------------------------------------------------------


                           RX                      TX


Port: internal-0/24
Total octets:              0                       0
Total packets:             0                       0
Unicast packets:           0                       0
Broadcast packets:         0                       0
Multicast packets:         0                       0
MAC control frames:        0                       0
CRC alignment errors:      0
Oversize packets:          0
Undersize packets:         0
```

```
Jabber packets:           0
Fragments received:       0
Ifout errors:             0
Packet drop events:       0
64        octets frames:  0
65-127    octets frames:  0
128-255   octets frames:  0
256-511   octets frames:  0
512-1023  octets frames:  0
1024-1518 octets frames:  0
Rate packets per second:  0                      0
Rate bytes per second:    0                      0

...

Port: vcp-0
Total octets:             0                      0
Total packets:            0                      0
Unicast packets:          0                      0
Broadcast packets:        0                      0
Multicast packets:        0                      0
MAC control frames:       0                      0
CRC alignment errors:     0
Oversize packets:         0
Undersize packets:        0
Jabber packets:           0
Fragments received:       0
Ifout errors:             0
Packet drop events:       0
64        octets frames:  0
65-127    octets frames:  0
128-255   octets frames:  0
256-511   octets frames:  0
512-1023  octets frames:  0
1024-1518 octets frames:  0
Rate packets per second:  0                      0
Rate bytes per second:    0                      0

Port: vcp-1
Total octets:             0                      0
Total packets:            0                      0
Unicast packets:          0                      0
Broadcast packets:        0                      0
```

```
Multicast packets:          0                       0
MAC control frames:         0                       0
CRC alignment errors:       0
Oversize packets:           0
Undersize packets:          0
Jabber packets:             0
Fragments received:         0
Ifout errors:               0
Packet drop events:         0
64        octets frames:    0
65-127    octets frames:    0
128-255   octets frames:    0
256-511   octets frames:    0
512-1023  octets frames:    0
1024-1518 octets frames:    0
Rate packets per second:    0                       0
Rate bytes per second:      0                       0

...

{master:0}
```

show virtual-chassis vc-port statistics member 0

```
user@switch>show virtual-chassis vc-port statistics member 0
fpc0:
--------------------------------------------------------------------------
Interface              Input Octets/Packets          Output Octets/Packets
internal-0/24          0           / 0               0           / 0
internal-0/25          0           / 0               0           / 0
internal-1/26          0           / 0               0           / 0
internal-1/27          0           / 0               0           / 0
vcp-0                  0           / 0               0           / 0
vcp-1                  0           / 0               0           / 0
internal-0/26          0           / 0               0           / 0
internal-0/27          0           / 0               0           / 0
internal-1/24          0           / 0               0           / 0
internal-1/25          0           / 0               0           / 0

{master:0}
```

Release Information

Command introduced in Junos OS Release 9.0.

The options `all-members`, `brief`, `detail`, `extensive`, and `local` were added in Junos OS Release 9.3 for EX Series switches.

RELATED DOCUMENTATION

clear virtual-chassis vc-port statistics

show virtual-chassis vc-port

Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis

7

CHAPTER

Knowledge Base