Virtual Chassis Feature Guide for Switches

Modified: 2018-12-27

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

Virtual Chassis Feature 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

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

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at https://www.juniper.net/documentation/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at https://www.juniper.net/books.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

   ```
   [edit]
   user@host# load merge /var/tmp/ex-script.conf
   load complete
   ```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   commit {
     file ex-script-snippet.xsl; }
   ```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see CLI Explorer.

Documentation Conventions

Table 1 on page xv defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xvi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** *metric***>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(***string1* **|** *string2* **|** *string3***)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [** *community-ids* **]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

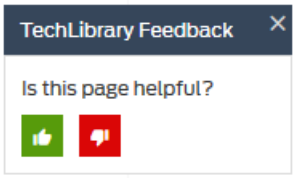

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://myjuniper.juniper.net

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://myjuniper.juniper.net.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://support.juniper.net/support/requesting-support/.

CHAPTER 1

Virtual Chassis Overview

Virtual Chassis Overview for Switches

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.

The following feature guides describe Virtual Chassis on different EX Series or QFX Series switches:

- *Virtual Chassis Feature Guide for EX2200, EX3300, EX4200, EX4500 and EX4550 Switches* covers configuring and maintaining EX2200, EX3300, EX4200, EX4500, and EX4550 Virtual Chassis.
- *Virtual Chassis Feature 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 Feature Guide for EX9200 Switches*.

- *Virtual Chassis Feature 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 Feature 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 master 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 master 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 62****.**

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 61 for details.

Virtual Chassis Configuration

You configure and manage nearly all aspects of an EX Series or QFX Series Virtual Chassis through the master 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 43 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 50 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 *master* 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 master Routing Engine role as the master 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 mastership election algorithm to select the member switches that assume the master 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 37
- Understanding How the Master in a Virtual Chassis Is Elected on page 57

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 master 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 58 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 master, 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 master 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.

Related Documentation

- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- Understanding Virtual Chassis Components on page 37
- Understanding Global Management of a Virtual Chassis on page 58
- Understanding How the Master in a Virtual Chassis Is Elected on page 57
- Understanding High Availability on an EX Series Virtual Chassis on page 68

Understanding EX Series Virtual Chassis

This topic introduces EX Series Virtual Chassis. An EX Series Virtual Chassis is a supported combination of interconnected EX2200, EX2300, EX3300, EX3400, EX4200, EX4500, EX4550, EX4300, or EX4600 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 Feature 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, EX2300 multigigabit model switches can also be combined with other EX2300 switches in the same Virtual Chassis, which operates as a non-mixed Virtual Chassis.
- EX3300 Virtual Chassis, composed of up to ten EX3300 switches.
- EX3400 Virtual Chassis, composed of up to ten EX3400 switches.
- 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.
- 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.
- 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 master 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 3 on page 26 lists the initial Junos OS release that supports each EX Series Virtual Chassis combination. "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 images might be different on different types of switches when mixed hardware models are supported together in a Virtual Chassis.

Table 3: Minimum Junos OS Release by Virtual Chassis Connection Type

Switch	EX2200 Switch	EX2300 Switch	EX3300 Switch	EX3400 Switch	EX4200 Switch	EX4300 Switch	EX4500 Switch	EX4550 Switch	EX4600 Switch
EX2200	12.2R1	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
EX2300	N/A	15.1X53-D50, 18.1R2 (MP models), 18.4R1 (MP and other models)	N/A	N/A	N/A	N/A	N/A	N/A	N/A
EX3300	N/A	N/A	11.3R1	N/A	N/A	N/A	N/A	N/A	N/A
EX3400	N/A	N/A	N/A	15.1X53-D50	N/A	N/A	N/A	N/A	N/A
EX4200	N/A	N/A	N/A	N/A	9.0R1	N/A	11.1R1	12.2R1	N/A
EX4300	N/A	N/A	N/A	N/A	N/A	13.2X50-D10, or 18.2R1 for MP models	N/A	N/A	13.2X51-D25, excludes EX4300 MP models
EX4500	N/A	N/A	N/A	N/A	11.1R1	N/A	11.1R1	12.2R1	N/A
EX4550	N/A	N/A	N/A	N/A	12.2R1	N/A	12.2R1	12.2R1	N/A

Table 3: Minimum Junos OS Release by Virtual Chassis Connection Type (continued)

Switch	EX2200 Switch	EX2300 Switch	EX3300 Switch	EX3400 Switch	EX4200 Switch	EX4300 Switch	EX4500 Switch	EX4550 Switch	EX4600 Switch
EX4600	N/A	N/A	N/A	N/A	N/A	13.2X51-D25, excludes EX4300 MP models	N/A	N/A	13.2X51-D25

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 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 50 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 the following VCP options:

- Network or uplink ports that can be configured into VCPs.
- Default-configured VCPs, which are configured as VCPs in the default factory configuration, but can alternatively be converted into network or uplink ports if desired and reconverted back into VCPs if needed.
- Dedicated VCPs, which can only be used 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, a Link Aggregation Group (LAG) bundle is formed if the links have the same speeds. For example, if you have four 40-Gbps links configured as VCPs between two member switches, a LAG with four member links at 160 Gbps of bandwidth is formed. Similarly, connecting two 10-Gbps links configured as VCPs between two member switches forms 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 be configured into and operate in one of three roles: master Routing Engine, backup Routing Engine, or linecard role. In some combinations of switches in a Virtual Chassis, certain switches are recommended or required to be in the Routing Engine roles. See "Understanding Virtual Chassis Components" on page 37 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 43), which automatically converts uplink ports into VCPs on the member switches on both sides of the new VCP links as they are cabled. This Virtual Chassis expansion method is also referred to as *autoprovisioning*.

EX2200 Switches in a Virtual Chassis

Up to four EX2200 and EX2200-C switches can be interconnected into an EX2200 Virtual Chassis, and cannot be mixed with any other type of switches in a Virtual Chassis.

You can configure and use any EX2200 1-Gbps 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 Gbps.

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).

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. EX2300 switches cannot be combined into a mixed Virtual Chassis with other EX Series or QFX Series switches, but can be interconnected into a non-mixed Virtual Chassis 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-Gbps Ethernet uplink ports as VCPs, and use those to interconnect the switches into a Virtual Chassis.

Interconnect 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 cannot use SFP transceivers on uplink ports to form an EX2300 Virtual Chassis.

In all EX2300 Virtual Chassis, any EX2300 switch model can be configured into any member switch role (master Routing Engine, backup Routing Engine, or linecard).

An EX2300 Virtual Chassis is configured and managed similarly to other EX Series and QFX Series Virtual Chassis. See the following for details on configuring or expanding an EX2300 Virtual Chassis:

- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis on page 94

EX3300 Switches in a Virtual Chassis

Up to ten EX3300 switches can be interconnected into a Virtual Chassis and cannot be mixed with any other type of switches.

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

An EX3300 Virtual Chassis is configured and managed similarly to 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)*

EX3400 Switches in a Virtual Chassis

You can connect up to ten of any models of EX3400 switches into a Virtual Chassis. EX3400 switches cannot be interconnected into 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. Uplink ports

with SFP tranceivers cannot be used as VCPs to interconnect EX3400 switches into a Virtual Chassis.

An EX3400 Virtual Chassis is configured and managed similarly to other EX Series and QFX Series Virtual Chassis. See the following for details on configuring or expanding an EX3400 Virtual Chassis:

- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis on page 94

EX4200, EX4500, and EX4550 Switches in a Virtual Chassis

Up to ten EX4200, EX4500, and EX4550 switches can be interconnected 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. Any SFP, SFP+, and XFP uplink ports on all of these switches can also be configured 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.

An EX4200, EX4500, or EX4550 Virtual Chassis is configured and managed similarly to 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 switcheds 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

Up to ten EX4300 switches can be connected together 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.

EX4300 multigigabit model switches and other EX4300 model switches can also be connected together into an EX4300 Virtual Chassis if the member switches are configured into mixed mode. In this case, the EX4300 member switches that are not multigigabit models must also be configured 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 EX4300 switches that are not multigigabit models can only be configured into linecard role. Multigigabit model EX4300 switches cannot interoperate in a mixed Virtual Chassis with any other types of switches.

EX4300 switches that are not multigigabit models can be included 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 50.

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-Gbps QSFP+ optical ports are configured as VCPs by default, and any 10-Gbps uplink module ports can also be configured 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-Gbps QSFP+ ports on the rear panel are dedicated VCPs. These are the only ports on EX4300 multigigabit model switches that can be used as VCPs, so any EX4300 Virtual Chassis that contains EX4300 multigigabit model switches can have only have VCP links that are 40 Gbps.

An EX4300 Virtual Chassis is configured and managed similarly to other EX Series and QFX Series Virtual Chassis. See the following for details on configuring and expanding an EX4300 Virtual Chassis:

- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis on page 94

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. Any EX4300 switches except multigigabit models (EX4300-48MP) can be interconnected 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 master and backup Routing Engine roles; EX4300 member switches cannot be configured into the Routing Engine role. As a result, at least two of the member switches in a mixed EX4600 Virtual Chassis must be EX4600 switches in the master and backup Routing Engine roles. See "Understanding Virtual Chassis Components" on page 37 for more information about Virtual Chassis member roles.

EX4600 switches do not have any ports that are configured into VCPs by default, but all 40-Gbps QSFP+ and 10-Gbps SFP+ optical ports on an EX4600 switch can be configured 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.

An EX4600 Virtual Chassis is configured, monitored, and maintained similarly to other EX and QFX Series Virtual Chassis. See the following for more details on configuring and expanding 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 81
- Adding an EX4600 Switch to a Mixed or Non-mixed Virtual Chassis on page 98

Release History Table

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1, EX2300 multigigabit model switches can also be combined with other EX2300 switches in the same Virtual Chassis, which operates as a non-mixed Virtual Chassis.

Related Documentation

- Virtual Chassis Overview for Switches on page 19
- Understanding Virtual Chassis Components on page 37
- Understanding Virtual Chassis Port Link Aggregation on page 61

Understanding QFX Series Virtual Chassis

This topic introduces QFX Series Virtual Chassis. A QFX Series Virtual Chassis is a supported combination of interconnected QFX3500, QFX3600, QFX5100, QFX5110, 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*.

This topic includes:

- Virtual Chassis Support on QFX Series Switches on page 33
- Basic Configuration of QFX Series Virtual Chassis on page 33
- QFX5200 Switches in a Virtual Chassis on page 34

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, and QFX5200 switches. EX4300 switches can also 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)
- 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 QFX Series switches.**

Basic Configuration of QFX Series Virtual Chassis

Some QFX Series switches can only form a Virtual Chassis with the same type of switches, while most 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 50 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:

- Any 40-Gbps QSFP+ ports on QFX5200 switches

NOTE: In Junos OS Release 17.3R2-S4, 100-Gbps QSFP28 ports are also supported as VCPs on QFX5200 switches.

- Any 100-Gbps or 40-Gbps QSFP28 ports on QFX5110 switches
- Any non-channelized 40-Gbps QSFP+ ports on QFX3500, QFX3600, QFX5100, or QFX5110 switches
- Any fixed 10-Gbps SFP+ ports on QFX Series switches that support these ports

QFX Series switches do not 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.

You can increase VCP bandwidth between member switches by configuring multiple interfaces between the same two switches into VCPs. When multiple VCPs interconnect the same two member switches, a Link Aggregation Group (LAG) or bundle is automatically formed when the VCPs are on interfaces supporting identical speeds. For example, if you have two 40-Gbps QSFP+ interfaces configured as VCPs between member switches, a LAG with two member links with 80Gbps of total bandwidth is formed. However, 10-Gbps SFP+ and 40-Gbps 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 be configured into and operate in one of three roles: master Routing Engine, backup Routing Engine, or linecard role. In some combinations of switches in a Virtual Chassis, certain switches are recommended or required to be in the Routing Engine roles. See "Understanding Virtual Chassis Components" on page 37 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 43), which automatically converts uplink ports into VCPs on the member switches on both sides of the new VCP links as they are cabled. This Virtual Chassis expansion method is also referred to as *autoprovisioning*.

QFX5200 Switches in a Virtual Chassis

Starting in Junos OS Release 17.3R2, Virtual Chassis is supported on QFX5200 switches.

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

QFX5110 Switches in a Virtual Chassis

Starting in Junos OS Release 17.3R1, Virtual Chassis is supported on QFX5110 switches.

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.

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, QFX5100-48T switches can be included in a QFX5110 Virtual Chassis.
- QFX5100-96S

QFX5100 Switches in a Virtual Chassis

Starting in Junos OS Release 13.2X51-D20, Virtual Chassis is supported on QFX5100 switches. 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: QFX5100-24Q switches should be configured 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 include 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 do not need to configure mixed mode. In a QFX5110 Virtual Chassis, we recommend to use QFX5110 switches in the master 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, QFX5100-48T switches can be included in a QFX5110 Virtual Chassis.
- QFX5100-96S

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. See "Upgrading a QFX5100 Switch with a USB Device to Join a QFX5110 Virtual Chassis or Virtual Chassis Fabric" on page 136.

QFX3500 and QFX3600 Switches in a Virtual Chassis

Virtual Chassis is supported on QFX3500 and QFX3600 series switches. To be included in a Virtual Chassis, QFX3500 and QFX3600 series switches must be configured as standalone switches and not as QFX node devices in a QFabric system.

QFX3500 and QFX3600 devices must be running a version of Junos OS for QFX devices that supports Virtual Chassis, and can be in a Virtual Chassis with up to ten total member switches composed of QFX3500 or QFX3600 switches only, or 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

Virtual Chassis is supported on EX4300 switches. Starting in Junos OS Release 13.2X51-D20, EX4300 switches except multigigabit models (EX4300-48MP) can be interconnected into a QFX Series Virtual Chassis with QFX3500 switches, QFX3600 switches, and QFX5100 switches.

A mixed or non-mixed QFX Series Virtual Chassis that includes EX4300 switches can contain up to ten total member switches including any combination of QFX3500, QFX3600, and QFX5100 switches with EX4300 switches.

Release History Table

Release	Description
17.3R2-S4	In Junos OS Release 17.3R2-S4, 100-Gbps QSFP28 ports are also supported as VCPs on QFX5200 switches.
17.3R2	Starting in Junos OS Release 17.3R2, Virtual Chassis is supported on QFX5200 switches.
17.3R2	Starting in Junos OS Release 17.3R2, QFX5100-48T switches can be included in a QFX5110 Virtual Chassis.
17.3R1	Starting in Junos OS Release 17.3R1, Virtual Chassis is supported on QFX5110 switches.
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, Virtual Chassis is supported on QFX5100 switches.
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 QFX Series Virtual Chassis with QFX3500 switches, QFX3600 switches, and QFX5100 switches.

Related Documentation

- Virtual Chassis Overview for Switches on page 19
- Understanding Virtual Chassis Components on page 37
- Configuring a QFX Series Virtual Chassis on page 86

Understanding Virtual Chassis Components

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: **For deployments with EX9200 switches, we recommend planning or moving 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.**

- A QFX Series Virtual Chassis is a supported combination of standalone QFX3500, QFX3600, QFX5100, QFX5110, or QFX5200 switches interconnected and managed as a single chassis. 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 discuss Virtual Chassis Fabric components. For information on Virtual Chassis Fabric components, 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 4 on page 39 lists the maximum member switch support by EX Series Virtual Chassis and Junos OS release.

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

Maximum Member Switch Support	Initial Junos OS Release
EX2200 Virtual Chassis	12.2R1—Initial release. Support for up to four EX2200 member switches.
EX2300 Virtual Chassis	15.1X53-D50—Initial release. Support for up to four EX2300 member switches. 18.1R2—Support for up to four multigigabit EX2300 (EX2300-24MP and EX2300-48MP) member switches. 18.4R1—Starting in Junos OS Release 18.4R1, up to four of any model EX2300 member switches (including multigigabit models and any other EX2300 switches) can be combined in the same Virtual Chassis.
EX3300 Virtual Chassis	11.3R1—Initial release. Support for up to six EX3300 member switches 12.2R1—Starting in Junos OS Release 12.2R1, an EX3300 Virtual Chassis can support up to ten EX3300 member switches.
EX3400 Virtual Chassis	15.1X53-D50—Initial release. Support for up to ten EX3400 member switches
EX4200 Virtual Chassis	9.0R1—Initial release. Support for up to ten EX4200 member switches
EX4300 Virtual Chassis	13.2X50-D10—Initial release. Support for up to ten EX4300 member switches 13.2X50-D20—Starting in Junos OS Release 13.2X50-D20, EX4300 switch support was added in a mixed QFX Series Virtual Chassis or in a VCF. 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 ten 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.
EX4500 Virtual Chassis	11.1R1—Initial release. Support for up to two EX4500 switches 11.4R1—Support for up to ten EX4500 member switches
EX4550 Virtual Chassis	12.2R1—Initial release. Support for up to ten EX4550 switches

Table 4: Maximum Member Switch Support for Virtual Chassis by Junos OS Release (continued)

Maximum Member Switch Support	Initial Junos OS Release
EX4600 Virtual Chassis	13.2X51-D25—Initial release. Support for up to ten EX4600 switches
Mixed EX4200 and EX4500 Virtual Chassis	11.1R1—Initial release. Support for up to two EX4500 switches and up to eight EX4200 switches 11.2R1—Support for up to nine EX4200 switches 11.4R1—Support for up to nine EX4500 switches
Mixed EX4200 and EX4550 Virtual Chassis	12.2R1—Initial release. Support for up to ten total EX4200 and EX4550 switches
Mixed EX4200, EX4500, and EX4550 Virtual Chassis	12.2R1—Initial release. Support for up to ten total EX4200, EX4500, and EX4550 switches
Mixed EX4300 and EX4600 Virtual Chassis	13.2X51-D25—Initial release. Support for up to ten total EX4300 and EX4600 switches. EX4600 switches 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. Support for up to ten total EX4500 and EX4550 switches
EX9200 Virtual Chassis	13.2R2—Initial release. Support for up to two EX9200 switches.

Maximum Switch Support in a QFX Series Virtual Chassis (Including Mixed Virtual Chassis with EX Series Switches)

In a QFX5200 Virtual Chassis, you can interconnect up to a maximum of 3 standalone QFX5200 switches as a non-mixed Virtual Chassis. (QFX5200 switches cannot be mixed with other types of switches in a Virtual Chassis.)

For all other QFX Series Virtual Chassis, you can interconnect up to 10 standalone switches in the following supported combinations:

- QFX5110 switches or a combination of QFX5110 switches and QFX5100 switches (a non-mixed Virtual Chassis)
- QFX5100 switches (a non-mixed Virtual Chassis)
- QFX5100 switches with any combination of QFX3500 and QFX3600 switches, and EX4300 switches excluding the multigigabit models (a mixed mode Virtual Chassis)

- QFX3500 switches, or QFX3600 swtiches, or any combination of QFX3500 switches and QFX3600 switches (a non-mixed Virtual Chassis)
- Any combination of QFX3500 and QFX3600 switches with EX4300 switches excluding the multigigabit models (a mixed mode Virtual Chassis)

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 QFX5100 Virtual Chassis.

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; these ports can only be used as VCPs and cannot be reconfigured as network ports. Dedicated VCPs allow you to interconnect switches without requiring any additional interface configuration.

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

Most switches have optical or uplink ports that can also be configured as 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, you must configure the VCPs. Also, 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 43).

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 what transceivers and cables are supported to 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 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 multigigabit EX2300 models)	None	None	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 uplink ports
EX3400	None	All QSFP+ uplink ports	Any SFP+ uplink ports NOTE: You cannot use ports with SFP transceivers as VCPs on EX3400 switches to form a Virtual Chassis.
EX4200	2 ports on 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
EX4300 Multigigabit Models (EX4300-48MP)	4 40-Gbps QSFP+ ports on rear panel	None	None
EX4500 and EX4550	Two ports on the Virtual Chassis module	None	Any SFP+ 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
QFX3500 and QFX3600	None	None	Any non-channelized 40-Gbps QSFP+ interfaces
QFX5100	None	None	Any non-channelized 40-Gbps QSFP+ interfaces
QFX5110	None	None	Any 100-Gbps or 40-Gbps QSFP28 ports Any non-channelized 40-Gbps QSFP+ interfaces

Table 5: VCP Options by Switch Type (continued)

Switch	Dedicated VCPs	Default VCPs	Ports that can be configured as VCPs
QFX5200	None	None	Any 40-Gbps QSFP+ ports In Junos OS Release 17.3R2-S4, 100-Gbps QSFP28 ports are also supported 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 referred to as *autoprovisioning* the new member.

Ports that are configured as VCPs by default on a switch or that were previously configured into VCPs must be converted back into network ports using the **request virtual-chassis vc-port delete** command for the port to be eligible for automatic VCP conversion. A port that has been automatically converted into a VCP is not automatically converted back into a network port when you remove a switch from a Virtual Chassis and disconnect the link.

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

- 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 occur even if the ports are manually converted into VCPs to create the redundant VCP link with a dedicated VCP link, so the reboot is required to avoid traffic looping in that case as well.)

- Starting in Junos OS Releases 14.1X53-D47, 17.4R2, 18.1R3, 18.2R2, and 18.3R1 for EX4300, EX4600, and QFX Series Virtual Chassis (which have the automatic VCP conversion feature enabled by default), you can choose to disable the feature by configuring the no-auto-conversion statement 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 interfaces between the same two switches into VCPs. When multiple VCPs interconnect the same two member switches, a Link Aggregation Group (LAG) or bundle is automatically formed when the VCPs are on interfaces supporting identical speeds. For example, if you have two 40-Gbps QSFP+ interfaces configured as VCPs between member switches, a LAG with two member links with 80-Gbps of total bandwidth is formed. However, 10-Gigabit SFP+ and 40-Gbps QSFP+ interfaces configured as VCPs 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 forwarded by the Virtual Chassis. See "Understanding Virtual Chassis Port Link Aggregation" on page 61 for details on the difference between VCP LAGs and network interface LAGs within a Virtual Chassis.

Master Routing Engine Role

In a Virtual Chassis, each member switch is assigned one of two roles: Routing Engine role or linecard role, and for Routing Engine role, receives a further designation as the master or backup Routing Engine.

The member that functions as the master in the Routing Engine role in the Virtual Chassis:

- Manages the member switches.
- Runs Junos OS for the switches as a master 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, one of the two members assigned as **routing-engine** functions as the master member. The selection of which member assigned as **routing-engine** functions as master and which as backup is determined by the software based on the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57.

In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm.

All switches that are not assigned the master or backup Routing Engine role function 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, any switch can be configured in any role in any configuration.
- In a mixed EX4300 Virtual Chassis composed of EX4300 multigigabit model (EX4300-48MP) and other EX4300 model switches, the switches in the Routing Engine role should always be EX4300 multigigabit model switches.
- In a mixed EX4600 Virtual Chassis with EX4300 switches, an EX4600 switch must assume the master 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. If the mixed Virtual Chassis does not contain QFX5100 switches, we recommend configuring QFX3500 or QFX3600 switches into the Routing Engine role.
- In a QFX5110 Virtual Chassis with QFX5110 and QFX5100 switches, we recommend configuring only QFX5110 switches into the Routing Engine role.

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 master Routing Engine role if the master fails.
- Runs Junos OS for the switches as a backup Routing Engine.
- Synchronizes with the master 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 master is unavailable.

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

In a preprovisioned configuration, one of the two members assigned as **routing-engine** functions in the backup role. The selection of which member assigned as **routing-engine** functions as master and which as backup is determined by the software based on the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57.

In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master 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, any switch can be configured in any role in any configuration.
- In a mixed EX4300 Virtual Chassis composed of EX4300 multigigabit model (EX4300-48MP) and other EX4300 model switches, the switches in the master and backup Routing Engine roles should always be EX4300 multigigabit model switches.
- In a mixed EX4600 Virtual Chassis with EX4300 switches, we strongly recommend configuring an EX4600 switch into the backup role. A mixed EX4600 and EX4300 Virtual Chassis must use an EX4600 member switch in the master role, and configuring an EX4600 switch into the backup role ensures that the Virtual Chassis remains up after a switchover event.
- 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. If the mixed Virtual Chassis does not contain QFX5100 switches, we recommend configuring QFX3500 or QFX3600 switches into the Routing Engine role.
- In a QFX5110 Virtual Chassis with QFX5110 and QFX5100 switches, we recommend configuring only QFX5110 switches into the Routing Engine role.

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 master.

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 makes it ineligible for functioning as a master or backup Routing Engine.

In a configuration that is not preprovisioned, the members that are not selected as master or backup function as linecard members of the Virtual Chassis configuration. The selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm. A switch with a mastership priority of 0 is always in the linecard role.

Any switch can function in the linecard role in a mixed or non-mixed Virtual Chassis.

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. If the mixed Virtual Chassis 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 one of those switches is powered on, it receives a member ID that can be seen by viewing the front-panel LCD or by entering the **show virtual-chassis** command. If the switch is powered on as a standalone switch, that member's member ID is always **0**. When the switch is interconnected with other switches in a Virtual Chassis configuration, its member ID is assigned by the master based on various factors, such as the order in which the switch was added to the Virtual Chassis configuration or the member ID assigned by a preprovisioned configuration. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57.

If the Virtual Chassis configuration previously included a member switch and that member was physically disconnected or removed from the Virtual Chassis configuration, its member ID is not available for assignment as part of the standard sequential assignment by the master. For example, you might have a Virtual Chassis configuration composed of member 0, member 2, and member 3, because member 1 was removed. When you add another member switch and power it on, the master assigns it as member 4.

The member ID distinguishes the member switches from one another. You use the member ID:

- To assign a mastership priority value to a member switch
- To configure interfaces for a member switch (The function is similar to that of a slot number on Juniper Networks routers.)
- To apply some operational commands to a member switch
- To display status or characteristics of a member switch

Mastership Priority

In a configuration that is not preprovisioned, you can designate the role (master or backup Routing Engine role, or linecard role) that a member switch assumes by configuring its mastership priority (from **0** through **255**). The mastership priority value is the factor in the master election algorithm with the highest precedence for selecting the master of the Virtual Chassis configuration. A switch with a mastership priority of **0** never assumes the backup or master Routing Engine role.

The default value for mastership priority is **128**. When a standalone switch is powered on, it receives the default mastership priority value. Because it is the only member of the Virtual Chassis configuration, it is also the master. When you interconnect a standalone switch to an existing Virtual Chassis configuration (which implicitly includes its own master), we recommend that you explicitly configure the mastership priority of the members that you want to function as the master and backup.

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

In a preprovisioned configuration, you assign the role of each member switch.

Virtual Chassis Identifier (VCID)

All members of a Virtual Chassis configuration share one Virtual Chassis identifier (VCID). This identifier is derived from internal parameters. When you are monitoring a Virtual Chassis configuration, the VCID is displayed in certain interface views and is also part of the **show virtual-chassis** output.

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 master 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 is shut down improperly, as follows:

- If the master is not available, the backup switch takes on the role of the master and its internal flash memory takes over as the alternate location for maintaining nonvolatile configuration memory.
- If a member switch is taken offline for repair, the master 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
18.4R1	Starting in Junos OS Release 18.4R1, up to four 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 ten 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	In Junos OS Release 17.3R2-S4, 100-Gbps QSFP28 ports are also supported 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, and QFX Series Virtual Chassis (which have the automatic VCP conversion feature enabled by default), you can choose to disable the feature by configuring the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level on the 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 QFX5100 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.
13.2X50-D20	Starting in Junos OS Release 13.2X50-D20, EX4300 switch support was added in a mixed QFX Series Virtual Chassis or in a VCF.
12.2R1	Starting in Junos OS Release 12.2R1, an EX3300 Virtual Chassis can support up to ten EX3300 member switches.

Related Documentation

- Virtual Chassis Overview for Switches on page 19
- *Understanding EX8200 Virtual Chassis Components*
- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- Understanding Mixed EX Series and QFX Series Virtual Chassis on page 50
- Configuring a QFX Series Virtual Chassis on page 86

- Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port on page 110
- Command Forwarding Usage with EX Series and QFX Series Virtual Chassis on page 121
- *Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet*
- *Example: Configuring an EX4500 Virtual Chassis with a Master 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

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 non-mixed 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.

- A QFX Series Virtual Chassis composed of only QFX3500 and QFX3600 switches
- A QFX5110 Virtual Chassis composed of QFX5110 and supported QFX5100 switches

QFX5200, EX3400, EX3300, EX2300, and EX2200 switches cannot be part of any mixed Virtual Chassis.

- Mixed and Non-mixed EX Series and QFX Series Virtual Chassis Summary on page 51
- Understanding the Routing Engine Role in a Virtual Chassis With Different Types of Switches on page 53
- Understanding QFX5100 and QFX5110 Switches in a Virtual Chassis on page 54
- Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Mixed Virtual Chassis on page 54
- Understanding Mixed EX4300 and EX4600 Virtual Chassis on page 55
- Understanding EX4300 Multigigabit and Other EX4300 Model Switches in a Mixed EX4300 Virtual Chassis on page 55
- Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis on page 56

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

Table 6 on page 52 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
	QFX5110	QFX5110 QFX5100 (with a "-qfx-5e-" Junos OS image)
	QFX5100	QFX5100
	QFX3600 QFX3500	QFX3600 QFX3500
	EX4600	EX4600
	EX4300 multigigabit models (EX4300-48MP) only	EX4300 multigigabit models (EX4300-48MP) only
	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 master Routing Engine role. We also recommend always configuring the same type of switch into the master and backup Routing Engine roles, to ensure that the switch operating as the master 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 master 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 master 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 master 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 master Routing Engine assumes the master Routing Engine role. Without user configuration, any switch might assume the master 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 master 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" on page 136.)

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 master 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 should not be configured into mixed mode.**

In a mixed EX4300 Virtual Chassis:

- You can interconnect up to ten member switches.
- EX4300 switches in the Virtual Chassis that are not multigigabit model switches must be configured with the **ieee-clause-82** option when setting mixed mode. See "Configuring an EX2300, EX3400, or EX4300 Virtual Chassis" on page 71 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.

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 master election process that decides member switch roles in a mixed Virtual Chassis is identical to the master election process in a non-mixed Virtual Chassis, so any member switch in a mixed Virtual Chassis can assume the master, 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

- Virtual Chassis Overview for Switches on page 19
- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- Understanding Virtual Chassis Components on page 37

Understanding How the Master in a Virtual Chassis Is Elected

NOTE: This topic does not apply to EX8200 Virtual Chassis.

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 *master* in a Routing Engine role and controls the Virtual Chassis configuration.

When a Virtual Chassis configuration boots, the Juniper Networks Junos operating system (Junos OS) on the switches automatically runs a master election algorithm to determine which member switch assumes the role of master.

The algorithm proceeds from the top condition downward until the stated condition is satisfied:

1. Choose the member with the highest user-configured mastership priority (255 is the highest possible value). A switch with a mastership priority of 0 will always stay in the linecard role.

2. Choose the member that was master the last time the Virtual Chassis configuration booted.

3. Choose the member that has been included in the Virtual Chassis configuration for the longest period of time. (For this to be a deciding factor, there has to be a minimum time lapse of 1 minute between the power-ons of the individual interconnected member switches.)

4. Choose the member with the lowest MAC address.

The variations among switches and switch models do not impact the master election algorithm.

To ensure that a specific member is elected as the master:

1. Power on only the switch that you want to configure as master of the Virtual Chassis configuration.

2. Configure the mastership priority of that member to have the highest possible value (255).

3. Continue to configure other members through the master member.

4. Power on the other members.

You can also specify the switch roles by preprovisioning your Virtual Chassis. Preprovisioning a Virtual Chassis allows you to manually assign the member ID and role for each switch in the 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 71, or "Configuring a QFX Series Virtual Chassis" on page 86.

Related Documentation

- Virtual Chassis Overview for Switches on page 19
- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- Understanding Virtual Chassis Components on page 37

Understanding Global Management of a Virtual Chassis

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, and it, therefore, has multiple console ports and multiple out-of-band management Ethernet ports located on the switches.

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 master switch, as shown in Figure 1 on page 59.

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

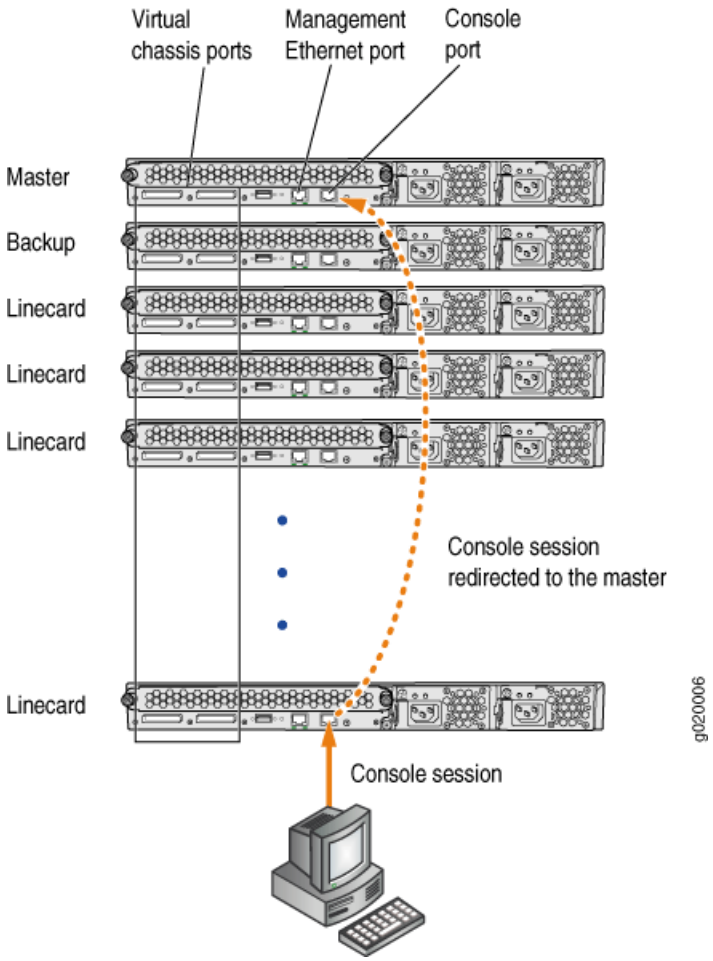

If the master becomes unavailable, the console session is disconnected from the old master and a new session is established with the newly elected master.

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 master member as shown in Figure 2 on page 60.

Figure 2: Management Ethernet Port Redirection to the VME Interface

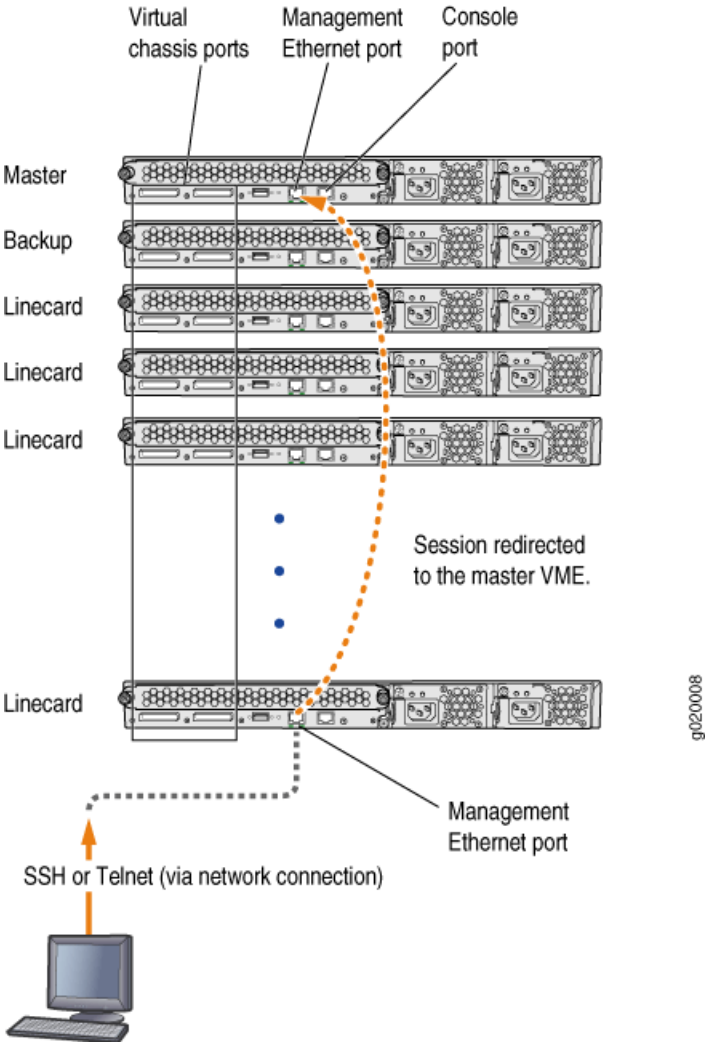

If the master management Ethernet link is unavailable, the session is redirected through the backup management Ethernet link. If there is no active management Ethernet link on the backup, the VME interface chooses a management Ethernet link on one of the linecard members, selecting the linecard member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

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.

Related Documentation

- Understanding Virtual Chassis Components on page 37
- *Example: Configuring an EX4200 Virtual Chassis with a Master 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

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 on page 61
- Virtual Chassis Port LAG Between Two Virtual Chassis Members on page 61

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.

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 110 or "Configuring a QFX Series Virtual Chassis" on page 86 for information about configuring uplink ports into VCPs.

Related Documentation

- Understanding EX Series Virtual Chassis on page 24
- *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 on page 32
- Configuring a QFX Series Virtual Chassis on page 86

Understanding Split and Merge in a Virtual Chassis

In a Virtual Chassis, two or more switches are connected together to form a unit that is managed as a single chassis. If there is a disruption to the Virtual Chassis configuration due to member switches failing or being removed from the configuration, the Virtual

Chassis configuration splits into two separate Virtual Chassis. This situation could cause disruptions in the network if the two separate configurations share common resources, such as global IP addresses. The split and merge feature provides a method to prevent the separate Virtual Chassis configurations from adversely affecting the network and also allows 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 but separate Virtual Chassis that have not previously been part of the same configuration into one Virtual Chassis configuration.

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

- What Happens When a Virtual Chassis Configuration Splits on page 63
- Merging Virtual Chassis Configurations on page 64

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 master election algorithm to select a new master for each of the two Virtual Chassis configurations. The new masters then determine whether their Virtual Chassis configuration remains active. One of the configurations remains active based on the following:

- It contains both the stable master and the stable backup (that is, the master and backup from the original Virtual Chassis configuration before the split).
- It contains the stable master 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 master 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: When you remove a member switch from a Virtual Chassis configuration, we recommend that you recycle the member ID using the request virtual-chassis recycle command.

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 mastership priority becomes the ID of the merged Virtual Chassis. The resulting merged Virtual Chassis configuration is active.

When you connect two Virtual Chassis configurations, the following events occur:

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 master election algorithm. The master election algorithm waits for the members to synchronize the topology information before running.

2. The master election algorithm merges the VCIDs of all the members.

3. Each member runs the master election algorithm to select a master and a backup from among all members with the same VCIDs. For more information, see "Understanding How the Master in a Virtual Chassis Is Elected" on page 57.

4. The master determines whether the Virtual Chassis configuration is active or inactive. (See "What Happens When a Virtual Chassis Configuration Splits" on page 63.)

5. If the Virtual Chassis configuration is active, the master assigns roles to all members. If the Virtual Chassis configuration is inactive, the master assigns all members the role of linecard.

6. When the other members receive their role from the master, they change their role to backup or linecard. They also use the active or inactive state information sent by the master to set their own state to active or inactive and to construct the Virtual Chassis member list from the information sent by the master.

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

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 master 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 on page 116

- Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge on page 119
- *Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge*
- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32

Understanding Automatic Software Update on Virtual Chassis Member Switches

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 on page 66
- Automatic Software Update Restrictions on page 66

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 master. When the master 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 master, the master 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 master 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" on page 136.

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 on page 135
- Configuring Automatic Software Update on Virtual Chassis Member Switches on page 116
- *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 master 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 master role remains a member of the Virtual Chassis. If the switch that was originally configured in the master role is removed from the Virtual Chassis, the MAC address of the current member switch in the master 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 Master of a Virtual Chassis on page 110
- Understanding EX Series Virtual Chassis on page 24
- *EX8200 Virtual Chassis Overview*
- Understanding QFX Series Virtual Chassis on page 32

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 master 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 69 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 69 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 on page 24

CHAPTER 2

Virtual Chassis Configuration

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis

This topic describes how to configure a non-mixed Virtual Chassis composed of EX2300, EX3400, or EX4300 switches, or how to configure a mixed EX4300 Virtual Chassis with EX4300 multigigabit model (EX4300-48MP) switches in combination with other EX4300 model switches, as follows:

- In Junos OS releases prior to 18.4R1, EX2300 and EX2300-C switches cannot be mixed with any other switches in a Virtual Chassis, and EX2300 multigigabit model switches, EX2300-24MP and EX2300-48MP, can only form a Virtual Chassis with other EX2300 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, with any of these switches in any role (master Routing Engine, backup Routing Engine, or linecard).

- EX3400 switches cannot be mixed with any other switches in a Virtual Chassis.
- EX4300 switches excluding multigigabit models can be configured into a non-mixed EX4300 Virtual Chassis, and EX4300 multigigabit model switches (EX4300-48MP) can also form a non-mixed EX4300 Virtual Chassis.

  In addition, EX4300 multigigabit model switches can be combined with other EX4300 model switches into a mixed EX4300 Virtual Chassis with the following parameters:

  - The Virtual Chassis must be configured into mixed mode.
  - The EX4300 switches that are not multigigabit models must be configured with a special port mode (**ieee-clause-82**) enabled on VCPs to interconnect with VCPs on multigigabit model member switches, which is set as an option when configuring these switches into mixed mode.
  - The switches in the Routing Engine role must be multigigabit model (EX4300-48MP) switches.

EX4300 switches excluding multigigabit models can be mixed in supported combinations with other switches in a Virtual Chassis or Virtual Chassis Fabric (VCF). Configuration procedures for these combinations are described in the following other configuration topics:

- Configuring a mixed Virtual Chassis that includes EX4600 and EX4300 member switches: "Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis" on page 81.
- Configuring a QFX Series Virtual Chassis that includes EX4300 member switches: "Configuring a QFX Series Virtual Chassis" on page 86.
- Configuring EX4300 switches as leaf nodes in a VCF: *Preprovisioning a Virtual Chassis Fabric* or *Autoprovisioning a Virtual Chassis Fabric*.

Use the following guidelines for planning the Virtual Chassis Port (VCP) connections in this configuration procedure:

- In a non-mixed EX4300 Virtual Chassis with only EX4300 multigigabit model (EX4300-48MP) switches, interconnect the member switches using the dedicated Virtual Chassis ports (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 dedicated VCPs on the multigigabit model switch members.
- In an EX2300 or EX3400 Virtual Chassis, or an EX4300 Virtual Chassis composed only of switches that are not multigigabit models, interconnect the member switches by using uplink ports configured as Virtual Chassis ports (VCPs).

  - EX2300 switches do not have any ports that are configured by default as VCPs.
  - The QSFP+ uplink ports on EX3400 and EX4300 switches, which support 40-Gbps speeds and can connect switches that are up to 492 feet (150 m) apart, are configured as VCPs by default, so you do not need to explicitly configure them.
  - On all of these switches, the SFP+ uplink ports, which support 10-Gbps speeds and can connect switches that are up to 6.2 miles (10 km) apart, can be configured into VCPs.
  - You cannot use uplink ports installed with SFP transceivers as VCPs.

- The simplest way to interconnect EX3400 or EX4300 switches into a non-mixed EX3400 or EX4300 Virtual Chassis is to install the switches within 492 feet of one another and interconnect them into a Virtual Chassis by using the QSFP+ ports. For an EX3400 or EX4300 Virtual Chassis, if the member switches must be installed in locations that are more than 492 feet apart (such as at a different site or at a distant location within the same site), or 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.
- 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 instance, if you have two QSFP+ ports and two SFP+ ports configured as VCPs connecting the same two member switches, the member switches would 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 master switch and the backup switch. 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 on both Routing Engines.

An EX2300, EX3400, or EX4300 Virtual Chassis can be configured with either of the following two options:

- A nonprovisioned configuration—The master sequentially assigns a member ID to other member switches. The role is determined by the mastership priority value and other factors in the master 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, with the exception of the interfaces on member switch 0, before the interfaces can send or receive traffic. The interfaces on member switch 0 are initially placed into the default VLAN; 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, or EX4300 Virtual Chassis with a Nonprovisioned Configuration File on page 74
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis with a Preprovisioned Configuration File on page 77

Configuring an EX2300, EX3400, or EX4300 Virtual Chassis with a Nonprovisioned Configuration File

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

To configure the Virtual Chassis using a nonprovisioned configuration:

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.

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 and EX4300 Virtual Chassis.

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

   

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

2. (Required for a mixed EX4300 Virtual Chassis only) Set the master switch into mixed mode, and reboot the switch to complete the configuration:

   ```
   user@device> request virtual-chassis mode 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 master 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 master switch apply to the entire Virtual Chassis configuration.

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

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

5. (Optional) Configure mastership 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 mastership priority value determines the roles in a non-provisioned Virtual Chassis configuration. The switches with the highest mastership priority values assume the master and backup roles. All other switches assume the linecard role.

If you do not configure the mastership priority for any switch in your Virtual Chassis, including when you do not configure the Virtual Chassis, all switches assume the default mastership priority of 128. The master election algorithm selects the roles for the member switches. In most cases, the switches that have been powered on the longest assume the master and backup roles when all Virtual Chassis member switches are configured with the same mastership priority. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57 for additional information on the master election algorithm.

A switch with a mastership priority of 0 never assumes the master or backup role.

NOTE: We recommend that you specify the same mastership priority value for the intended master and backup members.

6. (Optional: Recommended for a two-member Virtual Chassis) On the master 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 EX4300 Virtual Chassis only) Set each additional individual EX4300 multigigabit model (EX4300-48MP) switch into mixed mode and reboot the switch to complete the configuration:

```
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 to complete the configuration:

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

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

NOTE: 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.

NOTE: On EX3400 switches and on EX4300 member switches that are not multigigabit model switches, because QSFP+ ports are configured into VCPs by default, you do not usually have to perform this step when you are using a QSFP+ port as a VCP. You only need to configure a QSFP+ port as a VCP if you previously configured the QSFP+ port into a network port. If you previously configured the QSFP+ port into a network port, perform this step to configure the QSFP+ port back into a VCP.

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
```

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

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

Configuring an EX2300, EX3400, or EX4300 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 and EX4300 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 master or backup role. If you configure the member with a **line-card** role, it is not eligible to function in the master 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 master and backup Routing Engine roles must be EX4300 multigigabit model switches.

3. Power on only the switch that you plan to use as the master switch.

4. (Required for a mixed EX4300 Virtual Chassis only) Set the master switch into mixed mode, and reboot the switch to complete the configuration:

   ```
   user@device> request virtual-chassis mode 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 master 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 master switch apply to the entire Virtual Chassis configuration.

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

```
[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. 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 for a mixed EX4300 Virtual Chassis only) Set each additional individual EX4300 multigigabit model (EX4300-48MP) switch into mixed mode and reboot the switch to complete the configuration:

```
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 to complete the configuration:

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

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

NOTE: 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.

NOTE: On EX3400 switches and on EX4300 member switches that are not multigigabit model switches, because QSFP+ ports are configured into VCPs by default, you do not usually have to perform this step when you are using a QSFP+ port as a VCP. You only need to configure a QSFP+ port as a VCP if you previously configured the QSFP+ port into a network port. If you previously configured the QSFP+ port into a network port, perform this step to configure the QSFP+ port back into a VCP.

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
```

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

NOTE: You cannot modify the mastership priority when you are using a preprovisioned configuration. The mastership 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 mastership priority value. However, the member that was powered on first has higher prioritization according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57.

Related Documentation

- Configuring Mastership of a Virtual Chassis on page 107
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*

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

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.

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

In an EX4600 Virtual Chassis, you interconnect the EX4600 switches by using uplink ports configured as Virtual Chassis ports (VCPs). The 40-Gbps QSFP+ and 10-Gbps SFP+ uplink ports can be configured into VCPs. The uplink ports can be used to connect switches that are several miles apart, so you can connect switches in different buildings into the same Virtual Chassis.

You must configure QSFP+ or SFP+ uplink module ports into VCPs to create an EX4600 Virtual Chassis, or to interconnect EX4300 and EX4600 switches to form a mixed EX4300 and 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 master switch and the backup switch. 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 on both Routing Engines.

An EX4600 Virtual Chassis can be configured with either:

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

- Configuring an EX4600 Virtual Chassis with a Nonprovisioned Configuration File on page 82
- Configuring an EX4600 Virtual Chassis with a Preprovisioned Configuration File on page 84

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 master switch.

   If you are configuring a mixed EX4300 and EX4600 Virtual Chassis, you must use an EX4600 switch. A mixed EX4300 and EX4600 Virtual Chassis that uses an EX4300 in the master routing engine role is an unsupported configuration.

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

   ```
   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 an EX4600 Switch*.

   

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

4. (Optional) Configure the master 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 all other Virtual Chassis) Configure mastership priority for the member switches:

NOTE: If you are configuring a mixed EX4300 and EX4600 Virtual Chassis, configure the EX4600 switches with the highest mastership priorities to ensure EX4300 switches do not assume the master role. A mixed EX4300 and EX4600 switch operating with an EX4300 switch in the master role is an unsupported configuration.

```
[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 master 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 to complete the configuration:

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

9. 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 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 61.

NOTE: If you do not edit the Virtual Chassis configuration file, a nonprovisioned configuration is generated by default. The mastership priority value for each member switch is 128. The master role is selected by default. You can change the role that is performed by the members by modifying the mastership priority. See "Configuring Mastership of a Virtual Chassis" on page 107. We recommend that you specify the same mastership priority value for the desired master and backup members. In this example, the highest possible mastership priority has been assigned to two members. However, the member that was powered on first has higher prioritization according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57. The other members use the default mastership priority in this example, which configures them to function in the role of linecard.

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

Configuring an EX4600 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.

To configure 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 master or backup role. If you configure the member with a **line-card** role, it is not eligible to function in the master 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 master switch.

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

NOTE: You do not have to complete this step if you are configuring a Virtual Chassis that includes QFX3500 and QFX3600 switches only.

You must complete this step if your Virtual Chassis includes a mix of QFX5100, EX4300, and QFX3500 or QFX3600 switches.

```
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 an EX4600 Switch*.

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

6. (Optional) Configure the master 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
```

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 to complete the configuration:

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

12. 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 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 61.

NOTE: You cannot modify the mastership priority when you are using a preprovisioned configuration. The mastership 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 mastership priority value. However, the member that was powered on first has higher prioritization according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57.

Related Documentation

- Configuring Mastership of a Virtual Chassis on page 107
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*

Configuring a QFX Series Virtual Chassis

This topic discusses configuring a QFX Series Virtual Chassis. It does not apply to configuring a Virtual Chassis Fabric (VCF). For information on configuring a VCF, see *Understanding Virtual Chassis Fabric Configuration*.

You configure a QFX Series 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 32 and "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 50 for details on the switches that can be interconnected into a QFX Virtual Chassis, and the ports on those switches that can be used as VCPs.

A QFX Series Virtual Chassis can only be set up using QFX series switches configured in standalone mode.

- 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 QFX Series Virtual Chassis, you must download the software image for the standalone switch. QFX switches that are in a Junos Fusion cannot be part of a Virtual Chassis.

A QFX5110 Virtual Chassis can only be set up using QFX5110 and QFX5100 switches that are running the same Junos OS image, which must be an image that includes "-qfx-5e-" in the filename when the Junos OS software package is downloaded from the Software Center.

CAUTION: QFX5100 switches running a Junos OS image that includes "-qfx-5-" in the software package filename *must* be upgraded to the image filename that includes "-qfx-5e-" before being added 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" on page 136.

A QFX Series Virtual Chassis can be configured with either:

- A preprovisioned configuration—You can deterministically control the member ID and role assigned to a member switch by tying it to its serial number.
- A nonprovisioned configuration—The master sequentially assigns a member ID to other member switches. The role is determined by the mastership priority value and other factors in the master election algorithm.

NOTE: A Virtual Chassis configuration has two switches acting in the Routing Engine role—the master switch and the backup switch. 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 to both switches acting as Routing Engines.

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. Otherwise, as with other Junos OS releases, in Junos OS release 13.2X51-D20 you can interconnect up to ten of any other QFX5100 switches into a 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 QFX5100 Virtual Chassis.

NOTE: Feature licenses are required to configure advanced features on a Virtual Chassis.

For a Virtual Chassis deployment, two license keys are recommended for redundancy—one for the device in the master Routing Engine role and the other for the device in the backup Routing Engine role.

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

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

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 QFX Series Virtual Chassis on page 88
- Configuring a QFX Series Virtual Chassis with a Preprovisioned Configuration File on page 88
- Configuring a QFX Series Virtual Chassis with a Nonprovisioned Configuration File on page 91

Understanding the Licensing Requirements for a QFX Series Virtual Chassis

Feature licenses are required to configure advanced features on a Virtual Chassis.

For a Virtual Chassis deployment, two license keys are recommended for redundancy—one for the device in the master Routing Engine role and the other for the device in the backup Routing Engine role.

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

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

Configuring a QFX Series 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.

To configure a Virtual Chassis using a preprovisioned configuration:

NOTE: You can configure a QFX Series 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.

   

   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 master or backup Routing Engine . If you configure the member with a **line-card** role, it is not eligible to function as the master or backup. See "Understanding Virtual Chassis Components" on page 37 and "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 50 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 master switch.

4. Specify the identification parameters for the switch by completing the initial configuration. See *Configuring a QFX3500 Device as a Standalone Switch* or *Configuring a QFX3600 Device as a Standalone Switch*, *Configuring a QFX5100 Device*, *Configuring a QFX5110*, or *Performing the Initial Software Configuration for QFX5200 Switches*.

   

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

5. (Optional) Configure the master 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 master 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 50 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. 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 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 50 for details.

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

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

NOTE: SFP+, QSFP+, and 10Gbps copper links are automatically configured into VCPs when the preprovisioned configuration is set.

This step is, therefore, optional and should only be used when a VCP 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 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
```

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 61.

NOTE: You cannot modify the mastership priority when you are using a preprovisioned configuration. The mastership 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 mastership priority value. However, the member that was powered on first has higher prioritization according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57.

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 a QFX Series Virtual Chassis with a Nonprovisioned Configuration File

You can use nonprovisioned configuration to configure a QFX Series Virtual Chassis.

To configure the Virtual Chassis using a nonprovisioned configuration:

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

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

2. (Required for a mixed Virtual Chassis only) Set the master 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 50 for details.

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

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

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

4. (Optional) Configure the master 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 mastership priority for the other 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 master 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 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 50 for details.

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

9. 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 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
   ```

   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 61.

NOTE: If you do not edit the Virtual Chassis configuration file, a nonprovisioned configuration is generated by default. The mastership priority value for each member switch is 128. The master Routing Engine role is selected by default. You can change the role that is performed by the members by modifying the mastership priority. See "Configuring Mastership of a Virtual Chassis" on page 107. We recommend that you specify the same mastership priority value for the desired master and backup members. In this example, the highest possible mastership priority has been assigned to two members. However, the member that was powered on first has higher prioritization according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57. The other members use the default mastership priority in this example, which configures them to function in the role of linecard.

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

Release History Table

Release	Description
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 QFX5100 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.

Related Documentation

- Understanding QFX Series Virtual Chassis on page 32
- Understanding Mixed EX Series and QFX Series Virtual Chassis on page 50
- Configuring Mastership of a Virtual Chassis on page 107
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*

Adding a New Switch to an Existing EX2300, EX3400, or EX4300 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, which consists of either all EX4300 multigigabit model (EX4300-48MP) switches, or a combination of other EX4300 model switches excluding any multigigabit models.
- Add an EX4300 switch to a mixed EX4300 Virtual Chassis that consists of a combination of EX4300 multigigabit model switches and any other EX4300 switches.

NOTE: You cannot use this procedure to add EX4300 multigigabit model switches to an existing EX4300 Virtual Chassis that consists only of other non-multigigabit EX4300 model switches. EX4300 multigigabit model switches must be in the master 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 master 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.

See the following other references for details on how to add an EX4300 switch (excluding multigigabit models) to a mixed Virtual Chassis with EX4600 or QFX Series switches, or how to add an EX4300 switch (excluding multigigabit models) as a leaf node to a Virtual Chassis Fabric (VCF):

- 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 98.
- Adding an EX4300 switch to a mixed Virtual Chassis with QFX Series switches: "Adding a New Switch to an Existing QFX Series Virtual Chassis" on page 100
- Adding an EX4300 switch to a mixed VCF: *Adding a Device to a Virtual Chassis Fabric*.

NOTE: **To join an existing Virtual Chassis, new member switches must be running the same version of Junos OS that is running on the Virtual Chassis master. If you have configured the automatic software update feature in an existing Virtual Chassis, the master 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 66.

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 24 and "Understanding Virtual Chassis Components" on page 37 for details on the different EX Series switches, switch combinations, and switch roles that are supported or recommended in a Virtual Chassis.
- 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 master 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 43 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.

- (Optional) Considered deleting the **no-split-detection** configuration item if you are expanding a two-member Virtual Chassis with this option configured. The split detection and merge feature is recommended for Virtual Chassis configurations with more than two members, and is enabled by default when a Virtual Chassis is initially set up.

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 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 to complete the configuration.

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

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

   If the new switch is any other EX4300 model switch, you must also configure the switch with a special port mode using the **ieee-clause-82** option when configuring these switches into mixed mode, which enables any 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@device> request virtual-chassis mode mixed reboot all-members
   ```

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 interface as a Virtual Chassis Port (VCP), 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.
   - 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, or EX4300 Virtual Chassis on page 71
- Understanding Virtual Chassis Components on page 37

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. This procedure is also applicable when adding an EX4300 switch to an existing mixed EX4300 and EX4600 Virtual Chassis.

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 master 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 43 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+ interfaces as Virtual Chassis ports (VCPs), if needed:

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

   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. 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 on page 81
- Understanding Virtual Chassis Components on page 37

Adding a New Switch to an Existing QFX Series Virtual Chassis

This procedure applies to QFX Series Virtual Chassis. For the procedure on adding a switch to a Virtual Chassis Fabric (VCF), see *Adding a Device to a Virtual Chassis Fabric*.

You can use this procedure to add a device in a supported combination to an existing Virtual Chassis. A QFX Series Virtual Chassis can be a non-mixed or mixed Virtual Chassis consisting of a supported combination of switches interconnected using Virtual Chassis ports (VCPs). QFX series switches that can be members of a Virtual Chassis include QFX3500, QFX3600, QFX5100, QFX5110, or QFX5200 switches. EX4300 switches can also be members of a mixed QFX Series Virtual Chassis with QFX 3500, QFX3600, or QFX5100 switches. See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 50 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, be sure you have:

- Mounted the new switch in a rack.
- Enabled automatic software update on the Virtual Chassis. See "Configuring Automatic Software Update on Virtual Chassis Member Switches" on page 116.

CAUTION: When adding a QFX5100 switch that is installed with a "-qfx-5-" Junos OS image to a QFX5110 Virtual Chassis, you must first upgrade the QFX5100 switch before this step in the procedure so it is running a "-qfx-5e-" Junos OS image. 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" on page 136

- 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 master 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 43 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. 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*.

2. (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.

3. Interconnect the new switch to one member of the existing Virtual Chassis using an interface that can be configured into a VCP. See "Understanding Virtual Chassis Components" on page 37 for details on ports that can be used as VCPs.

   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.

4. Set the interconnecting interfaces for the new member switch as Virtual Chassis Ports (VCPs) if needed:

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

   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. 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 non-operational 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.

7. If further Virtual Chassis configuration is needed, see "Configuring a QFX Series Virtual Chassis" on page 86.

Related Documentation

- Configuring a QFX Series Virtual Chassis on page 86
- Understanding Virtual Chassis Components on page 37

Replacing a Member Switch of a Virtual Chassis Configuration

NOTE: This topic does not apply to Virtual Chassis Fabric (VCF). See *Removing a Device From a Virtual Chassis Fabric* for VCF information. This topic also does not apply to EX8200 Virtual Chassis.

You can replace a member switch in a Virtual Chassis without disrupting network service on the other members. You can retain the existing configuration of the member switch and apply it to a new member switch, or you can free up the member ID and make it available for assignment to a new member switch.

If you want to replace a member switch of a mixed Virtual Chassis that contains EX4200, EX4500, or EX4550 switches, see *Removing an EX4200, EX4500, or EX4550 Switch From a Mixed Virtual Chassis (CLI Procedure)*.

Otherwise, to replace a member switch of a Virtual Chassis, use the procedure that matches what you need to accomplish:

- Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch on page 103
- Remove, Repair, and Reinstall the Same Switch on page 103
- Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration on page 104

Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch

When you remove a member switch from the Virtual Chassis configuration, the master keeps that member switch's member ID in reserve. To make that member switch's member ID available for reassignment, issue the request virtual-chassis recycle command from the Virtual Chassis master.

NOTE: When you add or delete members in a Virtual Chassis configuration, internal routing changes might cause temporary traffic loss for a few seconds.

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 members. The master 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 master stores the configuration of the member that was removed. When you connect a different member switch, the master 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.

2. If the replacement member switch has been previously configured, revert that switch's configuration to the factory defaults. See *Reverting to the Default Factory Configuration for the EX Series Switch* for information about reverting to the factory default configuration on an EX Series switch or *Reverting to the Default Factory Configuration* for information about reverting to the factory default configuration on a QFX Series switch.

   The replacement member switch should be powered on and running with the factory default configuration at the end of this step.

3. (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" on page 116. 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 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: A QFX5110 Virtual Chassis can only be set up using QFX5110 and QFX5100 switches that are running the same Junos OS image, which must be an image that includes "-qfx-5e-" in the software package

filename when the Junos OS image is downloaded from the Software Center. If your replacement switch in a QFX5110 Virtual Chassis is a QFX5100 switch that was previously installed with and is running a Junos OS image from a software package filename that includes "-qfx-5-", 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" on page 136. The automatic software update feature cannot update a "-qfx-5-" image to a "-qfx-5e-" image.

4. 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. (The optical ports on the existing members in the Virtual Chassis where the replacement member will be connected should be configured as VCPs as well, 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 optical port VCPs for the switch after interconnecting a link into the existing Virtual Chassis. In either case, you must set the ports as VCPs that are used to interconnect the switch into the Virtual Chassis for the master 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" on page 110. For details on which ports can be configured as VCPs in a QFX Series Virtual Chassis, see "Understanding Virtual Chassis Components" on page 37.

5. 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, confirm the switch is part of the Virtual Chassis configuration by entering the **show virtual-chassis** command and viewing the output.

6. Cable the other VCP on the replacement member switch into the Virtual Chassis based on how you planned to interconnect the switch in Step 4.

   

   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.

7. To update the new member switch's current member ID to the member ID of the member switch that was removed from the Virtual Chassis configuration:

   - In a nonprovisioned Virtual Chassis, issue the request virtual-chassis renumber command on the master member switch.
   - In a preprovisioned Virtual Chassis, on the master 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 member'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 with a different member ID, associate the new member's serial number with the desired member ID and then delete the configuration item for the replaced member, 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
     ```

Related Documentation

- Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port on page 110
- *Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis (CLI Procedure)*
- Adding a New Switch to an Existing QFX Series Virtual Chassis on page 100
- *Monitoring the Virtual Chassis Status and Statistics on EX 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)*

Configuring Mastership of a Virtual Chassis

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 master 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 (master, 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 are required to be configured into the master or backup Routing Engine role, see "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 50.

NOTE: We recommend that you always use **commit synchronize** rather than **commit** to save configuration changes made for a Virtual Chassis, to ensure that the configuration changes are saved in both Routing Engines.

- Configuring Mastership Using a Preprovisioned Configuration File on page 108
- Configuring Mastership Using a Configuration File That Is Not Preprovisioned on page 108

Configuring Mastership Using a Preprovisioned Configuration File

To configure mastership using a preprovisioned configuration:

1. Note the serial numbers of the switches that you want to function as the master and backup Routing Engines.

   NOTE: Serial number values are case-sensitive.

2. Power on only the switch that you want to function as the master 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 master 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 mastership priority value when you are using a preprovisioned configuration. The mastership priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two members assigned the routing-engine role are assigned the same mastership priority value (128). However, the member that was powered on first has higher priority for the master election according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57. Only two members can be configured with the routing-engine role.

5. Specify the serial numbers of any other member switches that you are including in the Virtual Chassis configuration. You can also explicitly configure their role as **line-card**.

Configuring Mastership Using a Configuration File That Is Not Preprovisioned

To configure mastership of the Virtual Chassis through a configuration that is not preprovisioned:

1. Power on only the switch that you want to function as the master Routing Engine.

2. Configure the highest possible mastership priority value (**255**) for the member that you want to function as the master Routing Engine:

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

3. Configure the same mastership priority value (continue to edit the Virtual Chassis configuration on the master) for the member that you want to be the backup Routing Engine:

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

NOTE: We recommend that the master and backup have the same mastership priority value to prevent the master and backup status from switching back and forth between master and backup members in failover conditions.

4. Use the default mastership priority value (**128**) for the remaining member switches or configure the mastership priority to a value that is lower than the value specified for members functioning in the master and backup Routing Engine roles.

Related Documentation

- Understanding Mixed EX Series and QFX Series Virtual Chassis on page 50
- Configuring a QFX Series Virtual Chassis on page 86
- *Monitoring the Virtual Chassis Status and Statistics on EX 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)*

Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master 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 master switch as the system MAC base address. This process helps ensure a smooth transition of mastership with no disruption to network connectivity.

The MAC persistence timer is used in situations in which the master switch is no longer a member of the Virtual Chassis because it has been physically disconnected or removed. If the old master switch does not rejoin the Virtual Chassis before the timer elapses, the new master 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 67.

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 master switch is used as the system MAC base address; no MAC address changes occur within the Virtual Chassis even when the old master 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 a QFX Series Virtual Chassis on page 86
- Understanding Virtual Chassis Components on page 37

Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port

The procedure described in this topic can be used to set up Virtual Chassis ports (VCPs) to connect two switches together within the same EX Series or QFX Series Virtual Chassis. Switches that can be members of a Virtual Chassis might have:

- Uplink ports that can be configured into VCPs and converted back into uplink ports as needed.
- Default-configured VCPs—Ports that are already configured into VCPs with the default factory configuratio).
- Dedicated VCPs—Ports that can only be used as VCPs.

Under certain conditions, uplink ports that are supported as VCPs can automatically be converted into VCPs when interconnecting a member into an existing Virtual Chassis (see "Automatic Virtual Chassis Port (VCP) Conversion" on page 43). Otherwise, you must manually convert uplink ports into VCPs using this procedure. See "Virtual Chassis Port Options" on page 41 for a list of VCP options 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.

You typically configure an uplink port as a Virtual Chassis Port (VCP) for one of the following reasons:

- 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 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.

CAUTION: **If you configure an uplink 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 131 for more information.**

The following is a summary of VCP options or recommendations for switches that support Virtual Chassis:

- EX2200 switches—You explicitly configure VCPs to connect EX2200 switches together to form an EX2200 Virtual Chassis. See *Setting a Port on an EX2200 Switch as a Virtual Chassis Port (CLI Procedure)*.
- EX2300 switches—You explicitly configure 10-Gigabit uplink ports with SFP+ transceivers as VCPs to connect EX2300 switches or multigigabit EX2300 switches together to form an EX2300 Virtual Chassis. (You cannot use ports with SFP transceivers as VCPs on EX2300 switches.)
- EX3300 switches—Explicitly configuring VCPs is usually not needed when configuring an EX3300 Virtual Chassis. Uplink ports 2 and 3 on an EX3300 switch are configured as VCPs by default and, therefore, do not require user configuration to be set as VCPs. We recommend that you use this procedure to configure an uplink port on an EX3300 switch as a VCP only if you configured ports 2 and 3 as network uplink ports and the ports need to be reconfigured as VCPs, or when ports 2 and 3 cannot be used as VCPs for some reason. You can use this procedure to configure any uplink port on an EX3300 switch as a VCP.

- EX3400 switches—Explicitly configuring VCPs is usually not needed when configuring an EX3400 Virtual Chassis. The QSFP+ ports on EX3400 switches, which correspond to PIC (**pic-slot**) 1 ports 0 and 1, are configured as VCPs by default. Use this procedure only if you want to configure an SFP+ port on an EX3400 switch as a VCP, or if you want to reconfigure a QSFP+ port that had been configured into a network port back into a VCP. (You cannot use ports with SFP transceivers as VCPs on EX3400 switches.)
- EX4300 switches—Explicitly configuring VCPs is usually not needed when configuring an EX4300 Virtual Chassis. QSFP+ ports on EX4300 switches are configured as VCPs by default. Use this procedure only if you want to configure an SFP+ port on an EX4300 switch as a VCP, or if you want to reconfigure a QSFP+ port that had been configured into a network port back into a VCP.
- EX4200, EX4500, and EX4550 switches—EX4200, EX4500, and EX4550 switches have dedicated VCPs that are usually used to interconnect these switches into a Virtual Chassis. You can also interconnect EX4200, EX4500, and EX4550 switches that are beyond the reach of the dedicated Virtual Chassis cables as members of a Virtual Chassis by using the uplink ports—including the ports on the SFP uplink module, SFP+ uplink module, or XFP uplink module—and connecting the uplink ports. To use the uplink ports or SFP network ports for interconnecting member switches, you must explicitly set the uplink ports as VCPs.
- QFX Series switches—Explicitly setting up VCPs is required on any QFX Series switches to form a Virtual Chassis. QFX Series switches that support Virtual Chassis do not have any dedicated or default-configured VCPs, but you can configure any non-channelized QSFP+ ports (or QSFP28 ports, where supported) into VCPs.

When an uplink port is set as a VCP, it cannot be used for any other purpose. You can set one port as a VCP and configure other uplink ports in trunk mode as an uplink to another switch. You cannot use channelized ports as VCPs.

Before you set an uplink port as a VCP:

1. Verify that the port can be used as a VCP 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. Power on and connect to the switch that you plan to designate as the master of the Virtual Chassis.

   

   NOTE: **Do not power on the other switches at this point.**

4. Run EZSetup on the switch that you are configuring to be the master. Follow the prompts to specify the hostname and other identification, time zone, and network

properties. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* for details. The properties that you specify for the master apply to the entire Virtual Chassis, including all the member switches that you later interconnect with the master.

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 master 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 mastership of the Virtual Chassis by using either the nonprovisioned or preprovisioned configuration. See "Configuring Mastership of a Virtual Chassis" on page 107 for details.

NOTE: A Virtual Chassis has two Routing Engines, one in the master 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 a Virtual Chassis 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.

You can set uplink VCPs on a standalone switch before interconnecting one link into an existing Virtual Chassis, or set uplink VCPs for a switch after interconnecting one link into an existing Virtual Chassis. In either case, you must set the ports as VCPs that are used to interconnect the switch into the Virtual Chassis for the master to detect the switch and complete the process of adding it as a member.

NOTE: We recommend that you set up two uplink VCP connections within each wiring closet for redundancy.

1. Setting an Uplink VCP Between the Member Switches in a Virtual Chassis on page 114
2. Setting an Uplink VCP on a Standalone Switch on page 114

Setting an Uplink VCP Between the 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: If you use the SFP+ uplink module, you must configure all member switches to support either 1-gigabit SFP transceivers or 10-gigabit SFP+ transceivers on EX4200 switches. See *Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure)*.

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 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
   ```

2. Set one uplink port of member 1 as a VCP.

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

   This step includes the member *member-id* option, because it is executed on a different member switch than the local member switch.

   

   NOTE: You can also connect to a member switch individually using the request session member command, and set a VCP locally on that member. (As in Step 1, you do not specify the member option in that case.)

Setting an Uplink 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 master 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 the VCP uplink ports 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. If you have additional members in the second wiring closet, set a redundant VCP uplink on another member switch by issuing the **request virtual-chassis vc-port** command.

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)*
- *Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets*
- *Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File*
- Replacing a Member Switch of a Virtual Chassis Configuration on page 103
- *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 in a Virtual Chassis. You can disable the split and merge feature. If you disable the split and merge feature and the Virtual Chassis splits, both parts of the split Virtual Chassis configuration remain active.

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 remains inactive. If the Routing Engines end up in different parts of the split Virtual Chassis configuration and the rest of the member switches are configured as having linecard roles, then a backup Routing Engine might not be selected for either part.

We recommend disabling split and merge on a Virtual Chassis with two member switches. A two-member switch Virtual Chassis that has disabled split and merge can reform more quickly and with less complications as a result of the feature being disabled.

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 enable the split and merge feature again (recommended 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 on page 62

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 66.

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" on page 136.

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 on page 66
- Understanding Mixed EX Series and QFX Series Virtual Chassis on page 50

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 on page 62

Configuring Graceful Routing Engine Switchover in a Virtual Chassis

In a Virtual Chassis, one member switch is assigned the master role and has the master Routing Engine. Another member switch is assigned the backup role and has the backup Routing Engine. Graceful Routing Engine switchover (GRES) enables the master and backup Routing Engines in a Virtual Chassis configuration to switch from the master 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 master 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 mastership 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 Master and Backup in a Single Wiring Closet*
- *High Availability Features for EX Series Switches Overview*
- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32

CHAPTER 3

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 122 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 122 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**
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**

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 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**
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.	master 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**
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**

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 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**
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.	master 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**
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**

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 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 on page 19
- Understanding Virtual Chassis Components on page 37

Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member

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 master, 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 master 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 Mastership of a Virtual Chassis on page 107
- Configuring a QFX Series Virtual Chassis on page 86
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*

Verifying That Virtual Chassis Ports Are Operational

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

Purpose Verify that a Graceful Routing Engine switchover (GRES) between two member switches acting as the master and backup routing engines in a Virtual Chassis has occurred.

Action On the master 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 master 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 master and backup Routing Engines have switched roles:

NOTE: Member ID 1 is now the master 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 master and the master Routing Engine became the backup.

Related Documentation

- Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 119

Troubleshooting an EX Series Virtual Chassis

This topic describes the following troubleshooting issues for a Virtual Chassis:

A Disconnected Member Switch's ID Is Not Available for Reassignment

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 master 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

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

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

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

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)*

CHAPTER 4

Upgrading Software on a Virtual Chassis

Understanding Software Upgrades in a Virtual Chassis

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 a particular member 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. You might need to specify multiple Junos OS images when manually or automatically upgrading a mixed Virtual Chassis. For example, 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 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 on page 136
- Nonstop Software Upgrade on page 136

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 a Virtual Chassis. See "Understanding Automatic Software Update on Virtual Chassis Member Switches" on page 66 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 and the Virtual Chassis that support this feature, see:

- Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis on page 142
- Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade on page 150
- (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)*

Related Documentation

- Understanding Virtual Chassis Components on page 37
- Configuring Automatic Software Update on Virtual Chassis Member Switches on page 116
- *Installing Software on an EX Series Switch with a 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

Use this procedure to upgrade a standalone QFX5100 switch running "QFX 5 Series" Junos OS software to a "QFX 5e Series" software image that enables the QFX5100 to join a QFX5110 Virtual Chassis or Virtual Chassis Fabric (VCF).

A QFX5110 Virtual Chassis or VCF can be composed of a combination of QFX5110 and QFX5100 switches that must be running compatible Junos OS "QFX 5e Series" software. This procedure explains how to identify the compatible install package for QFX5100 switches, create a USB boot device as the installation media, and use the boot device to install the software on a QFX5100 switch. After installing the new software image, when you reboot the QFX5100, the switch is running a compatible software image and can be successfully added into a QFX5110 Virtual Chassis or VCF. Upgrading using a USB device is required in this case due to the differences in the boot structure and host OS software of the two types of devices.

NOTE: If a QFX5100 switch is already running a "QFX 5e Series" image (software package filename contains the string "-qfx-5e-"), you are not required to use this procedure to upgrade the switch before adding it to a QFX5110 Virtual Chassis or VCF. The Virtual Chassis or VCF can successfully update a QFX5100 member to the same version of Junos OS using the usual supported software update methods as needed during initial configuration or when adding or replacing members. See "Understanding Software Upgrades in a Virtual Chassis" on page 135 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 software labeled "QFX 5 Series" on the software download page, 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 software labeled "QFX 5e Series" (or simply "5e Series") on the software download page, and employ a secure-boot method at startup that is not used by QFX5100 members running a "-qfx-5-" or a "-qfx-5e-" software image. As a result, the corresponding install media and software package filenames for QFX5110 switches include the distinguishing strings "-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
```

Both QFX5110 and QFX5100 switches can run the *same* "-qfx-5e-" software images, and when inter-operating in a Virtual Chassis or VCF, they must be running the same "-qfx-5e-" software image that automatically either employs the secure-boot method (when booting on QFX5110 switches) or does not (when booting on QFX5100 switches). Due to platform and image differences, for a QFX5100 running a "-qfx-5-" image to successfully join a QFX5110 Virtual Chassis or VCF, you must initially upgrade the QFX5100 switch to run a "-qfx-5e-" image using a USB boot device *without* using the secure-boot method.

Install media packages with or without the secure-boot method that install the same "-qfx-5e-" software image for both switch types have the same filename, except the install media filename with the secure-boot method includes the "-secure-" keyword, and the install media filename for without the secure-boot method does not. For example, the following install media packages install the same software image for QFX5110 and QFX5100 switches, respectively, but only the first install media package uses the secure boot method during the installation:

Install media package for QFX5110 members (with the secure-boot method):

```
install-media-host-usb-qfx-5e-x86-64-17.3R1.6-secure-signed.tgz
```

Compatible install media package for QFX5100 switches (without using the secure-boot method):

```
install-media-host-usb-qfx-5e-x86-64-17.3R1.6-signed.tgz
```

As a result, when creating the USB boot device to upgrade a QFX5100, you download the install media filename *without* the "secure" keyword that otherwise matches the software package name for the image running on the QFX5110 switches in your Virtual Chassis or VCF. For example:

```
install-media-host-usb-qfx-5e-x86-64-17.3R1.6-signed.tgz
```

After a QFX5100 switch is upgraded to a "-qfx-5e-" image using this procedure, the same software image updates can be installed on either type of switch in the Virtual Chassis or VCF from the corresponding "jinstall-host-qfx-5e-" install package file *with* the "secure" keyword. When the image boots, it determines whether or not to use the secure-boot method based on the type of switch on which it is running. The QFX5110 Virtual Chassis or VCF can also successfully update the member, if necessary, using the automatic software update feature during configuration or when adding or replacing members.

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-" software image to run as a standalone switch, you similarly 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 upgrade a QFX5100 switch to run the corresponding 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 used 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 QFX5100 switches must initially be upgraded without using the same secure-boot method employed by QFX5110 switches.

NOTE: The Junos OS software running on the QFX5110 members must be the "-qfx-5e-" 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-" software image on a QFX5100 switch, the same software image that supports secure boot runs on either type of switch, and determines the appropriate boot method to use based on the type of switch on which it is running. As a result, for subsequent 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 137) 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-" 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

A standalone QFX5100 switch running "QFX 5 Series" software must be upgraded to "QFX 5e Series" software before the switch can join a QFX5110 Virtual Chassis or VCF. The upgrade requires booting and installing 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 139. 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 proceed as follows:

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**. 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.

   When installation is complete, 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-" 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).

Related Documentation

- Understanding Mixed EX Series and QFX Series Virtual Chassis on page 50
- *Understanding Mixed Virtual Chassis Fabric*
- Configuring a QFX Series Virtual Chassis on page 86
- *Understanding Virtual Chassis Fabric Configuration*

Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis

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 introduces NSSU on EX Series and QFX Series Virtual Chassis that support this feature, except 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)*. For information on using NSSU with a Virtual Chassis Fabric (VCF), see *Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric*.

NOTE: Because NSSU upgrades the software on each Virtual Chassis and mixed Virtual Chassis member one at a time, an upgrade 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. The members of a Virtual Chassis or VCF in an upgrade group are upgraded simultaneously, reducing the amount of time it takes to complete an upgrade. See "Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade" on page 147.

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, interface, kernel, and routing protocol information is preserved.
- Minimal disruption to network traffic—NSSU minimizes network traffic disruption by upgrading member switches one at a time, enabling the master and backup members to maintain their master and backup roles (although mastership will change) without disruption to traffic, and permitting traffic to continue to flow through members in the linecard role that are not being upgraded.

Requirements for Performing an NSSU

The following are requirements for performing an NSSU for a Virtual Chassis or mixed Virtual Chassis:

- All Virtual Chassis or mixed Virtual Chassis members and all Routing Engines must be running the same Junos OS release.
- Graceful Routing Engine switchover (GRES) must be enabled.
- Nonstop active routing (NSR) must be enabled for applicable platforms.

  Although nonstop bridging (NSB) is not required to perform an NSSU, we recommend enabling NSB before performing an NSSU on applicable platforms. Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU. In releases prior to Junos OS Release 16.1, 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 or mixed 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 an EX Series switch in a mixed Virtual Chassis is upgraded 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.

The following are requirements for the Virtual Chassis or mixed Virtual Chassis members being upgraded using NSSU:

- The Virtual Chassis or mixed Virtual Chassis members 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 Virtual Chassis or mixed Virtual Chassis master and backup must be adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in linecard roles are rebooting.
- The Virtual Chassis or mixed Virtual Chassis must be preprovisioned so that the linecard role has been explicitly assigned to member switches acting in a linecard role. During an NSSU, the Virtual Chassis and mixed Virtual Chassis members must maintain their roles—the master and backup must maintain their master and backup roles (although mastership will change), and the remaining switches must maintain their linecard roles.
- A two-member Virtual Chassis or mixed Virtual Chassis must have **no-split-detection** configured so that the Virtual Chassis or mixed Virtual Chassis does not split when an NSSU upgrades a member.

NOTE: For the EX4300 Virtual Chassis, you should enable the **vcp-no-hold-time** statement at the **[edit virtual-chassis]** hierarchy level before performing a software upgrade using NSSU. If you do not enable the **vcp-no-hold-time** statement, 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 62

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 or mixed Virtual Chassis master verifies that:
   - The backup is online and running the same software version.
   - Graceful Routing Engine switchover (GRES) is enabled, and, if applicable, nonstop active routing (NSR) is enabled.
   - The Virtual Chassis or mixed Virtual Chassis member has a preprovisioned configuration.
2. The master transfers 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 master 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 number of parallel copy operations is determined by a default algorithm 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 aborts 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, after an NSSU copy of the new software image to a member fails, the master performs an additional error recovery measure to remove the new software from the members to which it was already transferred.

3. The master restarts the backup with the new software, and the backup resynchronizes with the master.

4. The master loads and reboots member switches that are in the linecard role, one at a time. The master 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 or VCF 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.
   - 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 aborts 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. This action cleanly brings up all members at the same time running the new software, which recovers stable Virtual Chassis operation more quickly than having an unstable Virtual Chassis running different versions of the software trying to converge.

6. When all members that are in the linecard role have been upgraded, the master performs a graceful Routing Engine switchover, and the upgraded backup becomes the master.
7. The software on the original master is upgraded and the original master is automatically rebooted. After the original master has rejoined the Virtual Chassis, you can optionally return control to it by requesting a graceful Routing Engine switchover.

NSSU Limitations

You cannot use an 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 cannot 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 do so by rebooting 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

A Virtual Chassis or mixed Virtual Chassis must be running a Junos OS release that supports NSSU before you can perform an NSSU. If a Virtual Chassis or mixed Virtual Chassis is running a software version that does not support NSSU, use the **request system software add** command. See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 50 for information on which switches are supported on a mixed Virtual Chassis.

Overview of NSSU Configuration and Operation

You must ensure that the configuration of the switch or Virtual Chassis meets the requirements described in "Requirements for Performing an NSSU" on page 142. NSSU requires no additional configuration.

You perform an NSSU by executing the **request system software nonstop-upgrade** command. For detailed instructions on how to perform an NSSU, see the topics in Related Documentation.

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 master 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, after an NSSU copy of the new software image to a member fails, the master 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 on page 150
- Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 119
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

Nonstop software upgrade (NSSU) enables you to upgrade software using a single command and with minimal disruption to network traffic on supporting switches. To reduce the total time required to complete an NSSU operation on all switches being upgraded, you can configure line-card upgrade groups on the following supported platforms:

- an EX6200 or EX8200 switch with redundant Routing Engines
- an EX8200 Virtual Chassis
- a QFX3500, QFX3600, and QFX5100 Virtual Chassis
- a Virtual Chassis Fabric

In its default configuration, NSSU upgrades each line card in a switch or 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 allows you to minimize disruption to traffic if you have configured 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 define an upgrade group for NSSU, NSSU upgrades the line cards or members in the 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, 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)*.

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**.

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]
```

To configure line-card upgrade groups on a QFX Series Virtual Chassis or mixed Virtual Chassis, or a VCF:

NOTE: For Virtual Chassis or VCFs comprised of fixed-chassis switches that do not have separate line cards, you use the **upgrade-group** configuration statement with the **fpcs** option to specify the Virtual Chassis or VCF member IDs that you want to include in an upgrade group. The **member** hierarchy of the **upgrade-group** statement is not used.

- To create an upgrade group and add a Virtual Chassis or VCF member to the upgrade group:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs value
```

  For example, to create an upgrade group called **vcf** and add a line card member:

```
[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:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs [list-of-slot-numbers]
```

  For example, to create an upgrade group called **vcf** that contains members 1 and 2:

```
[edit chassis]
user@switch# set nssu upgrade-group vcf fpcs [1 2]
```

Related Documentation

- *Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches*
- *Understanding Nonstop Software Upgrade on EX Series Switches*
- Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis on page 142
- *Understanding Nonstop Software Upgrade on EX Series Switches*
- *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

You can use nonstop software upgrade (NSSU) to upgrade the software running on all member switches of a Virtual Chassis with minimal traffic disruption during the upgrade.

Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- Ensure that the Virtual Chassis or mixed Virtual Chassis is configured correctly to support NSSU. Verify that:
  - The Virtual Chassis or mixed Virtual Chassis members are connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
  - The Virtual Chassis or mixed Virtual Chassis master and backup are adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in linecard roles are rebooting.
  - The Virtual Chassis or mixed Virtual Chassis is preprovisioned so that the linecard role has been explicitly assigned to member switches acting in the linecard role. During an NSSU, the Virtual Chassis or mixed Virtual Chassis members must maintain their roles—the master and backup must maintain their master and backup roles (although mastership will change), and the other member switches must maintain their linecard roles.
  - A two-member Virtual Chassis has no-split-detection configured so that the Virtual Chassis does not split when an NSSU upgrades a member.
- 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.

  For applicable platforms, ensure that nonstop active routing (NSR) is enabled, and in that case graceful Routing Engine switchover will also be enabled—you need to check

only the state of nonstop active routing with the *show task replication* command to verify both NSR and GRES are enabled. See *Configuring Nonstop Active Routing on Switches* for more information.

- For EX4300 Virtual Chassis, you should enable the **vcp-no-hold-time** statement at the [**edit virtual-chassis**] hierarchy level 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 62.
- For QFX5100 Virtual Chassis, you should enable the **lc-reboot-delay** statement to configure a delay for when adjacent members in a line card group reboot. Without this option, when the next member reboots, after approximately two minutes after the previous member reboots and joins the Virtual Chassis, the previous rebooted member might not be ready to carry traffic at that time. This delay is needed when there are two adjacent line card members that have interfaces that are part of a common lag, otherwise traffic might drop.

  To configure the delay, issue the **set chassis nssu lc-reboot-delay *seconds* command**. We recommend that you configure 200s. The range of seconds is between 0 and 600.
- (Optional for applicable platforms) Enable nonstop bridging (NSB). Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on each member to an external storage device with 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. Because a graceful Routing Engine switchover occurs during the upgrade, the original Virtual Chassis or mixed Virtual Chassis backup is the new master.

During NSSU, the master 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 aborts 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 aborts due to a copy error, the new image is removed 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, performing 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 is not supported.**

To upgrade all members in a Virtual Chassis using NSSU:

1. Download the software package by following the procedure in the “Downloading Software Files with a Browser” section in *Installing Software Packages on QFX Series Devices*. If you are upgrading the software running on a mixed Virtual Chassis, download the software packages for the applicable switch types.
2. Copy the software package or packages to the Virtual Chassis or mixed Virtual Chassis. We recommend that you copy the file or files to the **/var/tmp** directory on the master.
3. Log in to the Virtual Chassis or mixed Virtual Chassis using the console connection or the virtual management Ethernet (VME) interface. Using a console connection allows you to monitor the progress of the master switch reboot.
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 reboot of the original master switch completes. To verify that the software on all Routing Engines in the Virtual Chassis or mixed Virtual Chassis members has been upgraded, enter the following command:

   ```
   user@switch> show version
   ```

6. To ensure that 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
   ```

   Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Related Documentation

- Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis on page 142
- *Configuring Dual-Root Partitions*

CHAPTER 5

Configuration Statements

aliases (Virtual Chassis)

Syntax

```
aliases {
  serial-number serial-number {
    alias-name alias-name;
  }
}
```

Hierarchy Level

[edit virtual-chassis]

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.

Description

Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.

An alias is not specified for a device until the alias name is specified using the **alias-name** keyword.

The alias appears in the **Alias-Name** field in the **show virtual-chassis** command.

Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring a QFX Series Virtual Chassis on page 86
- *Understanding Virtual Chassis Fabric Components*
- Understanding Virtual Chassis Components on page 37

alias-name (Virtual Chassis aliases)

Syntax alias-name *alias-name*;

Hierarchy Level [edit virtual-chassis aliases serial-number *serial-number*]

Release Information Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.

Description Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.

The alias appears in the **Alias-Name** field in the **show virtual-chassis** command.

Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.

In the following example, the **dc-floor-1** alias name is assigned to the member switch with the serial number AB0123456789.

set serial-number

```
[edit virtual-chassis aliases]
user@switch# set serial-number AB0123456789 alias-name dc-floor-1
```

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
<additional output removed for brevity>
```

Options ***alias-name***—The text label, or alias, assigned to the member switch by the user.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring a QFX Series Virtual Chassis on page 86

- *Understanding Virtual Chassis Fabric Components*
- Understanding Virtual Chassis Components on page 37

auto-sw-update

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]

Release Information

Statement 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.
Statement 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.
The **ex-4600** option introduced in Junos OS Release 13.2X51-D25.

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 the 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 same 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" on page 136.

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 a path to a Junos OS software image.

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.

Related Documentation

- *Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches*
- Configuring Automatic Software Update on Virtual Chassis Member Switches on page 116
- Understanding Software Upgrades in a Virtual Chassis on page 135
- *Understanding Software Upgrades in a Virtual Chassis Fabric*

fpcs (NSSU Upgrade Groups)

Syntax

```
fpcs (slot-number | [list-of-slot-numbers]);
```

Hierarchy Level

```
[edit chassis (EX Series) nssu upgrade-group group-name],
[edit chassis (EX Series) nssu upgrade-group group-name member member-id]
```

Release Information

Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.

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; QFX3500, QFX3600, and QFX5100 Virtual Chassis; or a Virtual Chassis Fabric (VCF).

For switches that have separate line cards, use this statement to assign one or more line cards to an NSSU upgrade group by specifying their slot numbers with this statement.

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: You do not use this statement with the member keyword in this case. When to use the member statement hierarchy is explained next.

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 with the member statement hierarchy 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.

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 on page 147

id

Syntax

id *id*;

Hierarchy Level

[edit virtual-chassis]

Release Information

Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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.

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 on page 119
- Configuring a QFX Series Virtual Chassis on page 86
- *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*

lag-hash

Syntax lag-hash (packet-based | source-port-based);

Hierarchy Level [edit virtual-chassis vc-port]

Release Information Statement introduced in Junos OS Release 12.1 for EX Series switches.

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

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.

location (Virtual Chassis)

Syntax

location *location*;

Hierarchy Level

[edit virtual-chassis member *member-id*]

Release Information

Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring a QFX Series Virtual Chassis on page 86
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis on page 81

mac-persistence-timer

Syntax mac-persistence-timer (*minutes* | **disable**);

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
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.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

Description Specify how long the Virtual Chassis or VCF continues to use the MAC address of the switch that was originally configured in the master role as the system MAC base address after the original master 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 master 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 master before using its own system MAC base address after the switch in the master 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 master 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.

Related Documentation

- Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis on page 110
- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*

mastership-priority

Syntax mastership-priority *number*;

Hierarchy Level [edit virtual-chassis member *member-id*]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Mastership priority option **0** introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

Description The mastership priority value is the most important factor in determining the role of the member switch within a nonprovisioned Virtual Chassis or VCF configuration. Other factors (see "Understanding How the Master in a Virtual Chassis Is Elected" on page 57) also affect the election of the master.

The mastership priority value takes the highest precedence in the master election algorithm. The member switch with highest mastership priority assumes the master Routing Engine role of the Virtual Chassis or VCF. Toggling back and forth between master and backup status in failover conditions is undesirable, so we recommend that you assign the same mastership priority value to both the master and the backup. Secondary factors in the master election algorithm determine which of these two members (that is, the two members that are assigned the highest mastership priority value) functions as the master of the Virtual Chassis or VCF.

This statement is not used for the EX8200 Virtual Chassis, which determines mastership by external Routing Engine uptime.

A switch with a mastership priority of **0** never takes the master or backup role.

Default 128

Options *number*—Mastership 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.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring a QFX Series Virtual Chassis on page 86
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis on page 81

member

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]
```

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 master 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.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*

- *Adding a Device to a Virtual Chassis Fabric*
- Configuring a QFX Series Virtual Chassis on page 86
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis on page 81
- *Configuring an EX9200 Virtual Chassis*
- Configuring a QFX Series Virtual Chassis on page 86

member (NSSU Upgrade Groups)

Syntax

```
member member-id {
  fpcs (slot-number | [list-of-slot-numbers]);
}
```

Hierarchy Level

[edit chassis (EX Series) nssu upgrade-group *group-name*]

Release Information

Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.

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.

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 on page 147

no-auto-conversion

Syntax no-auto-conversion;

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced in Junos OS Releases 14.1X53-D47, 17.4R2, 18.1R3, 18.2R2, and 18.3R1 for EX4300, EX4600, and QFX Series switches.

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 for EX Series and QFX Series switches in a Virtual Chassis that have automatic VCP conversion enabled by default, which includes EX4300, EX4600, and QFX Series switches that support Virtual Chassis. See "Automatic Virtual Chassis Port (VCP) Conversion" on page 43 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.

Related Documentation
- Understanding Virtual Chassis Components on page 37
- Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port on page 110
- Adding a New Switch to an Existing EX2300, EX3400, or EX4300 Virtual Chassis on page 94
- Adding an EX4600 Switch to a Mixed or Non-mixed Virtual Chassis on page 98
- Adding a New Switch to an Existing QFX Series Virtual Chassis on page 100

no-management-vlan

Syntax
```
no-management-vlan;
```

Hierarchy Level
[edit virtual-chassis member *member-id*]

Release Information
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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.

Related Documentation
- Understanding Global Management of a Virtual Chassis on page 58
- *Understanding Virtual Chassis Fabric Configuration*

no-split-detection

Syntax

```
no-split-detection;
```

Hierarchy Level

[edit virtual-chassis]

Release Information

Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

Description

Disable the split and merge feature in a Virtual Chassis or VCF configuration.

We recommend using this statement to disable the split and merge feature when configuring a two-member Virtual Chassis. Enabling this statement on a two-member Virtual Chassis ensures that both switches remain in the correct Virtual Chassis roles in the event of a Virtual Chassis split.

The split and merge feature is enabled by default when a Virtual Chassis is initially set up. If you set this option in a two-member Virtual Chassis and then expand the Virtual Chassis to add more members, we recommend deleting this 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.

Related Documentation

- *Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge*
- Disabling Split and Merge in a Virtual Chassis on page 116
- Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge on page 119
- *Disabling Split Detection in a Virtual Chassis Configuration*
- Understanding Split and Merge in a Virtual Chassis on page 62

nssu

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 (EX Series)]

Release Information

Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.
rcp-count statement introduced in Junos OS Release 14.1X53-D40 for QFX5100 switches.

Description

Configure optional parameters used in the nonstop software upgrade (NSSU) process.

NOTE: The request system software nonstop-upgrade command is used to initiate NSSU.

For the **rcp-count** statement: (QFX5100 Virtual Chassis and Virtual Chassis Fabric (VCF) only) Configure the number of parallel **rcp** sessions NSSU uses to copy the new software to multiple Virtual Chassis or VCF member switches at a time. See rcp-count for details.

For **upgrade-group** statements: Define a line-card upgrade group for NSSU, for switch configurations that support upgrade groups. All line cards or Virtual Chassis or VCF members in an upgrade group are upgraded to the new software version at the same time when an NSSU is initiated and at least one upgrade group is configured. Line-card upgrade groups are not required to initiate an NSSU, and are not supported on some EX Series switches or EX Virtual Chassis that support NSSU. See upgrade-group for details.

The remaining statements are explained separately. See CLI Explorer.

Default

If **rcp-count** is not configured, 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 no line-card upgrade groups are defined, NSSU upgrades line cards and 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.

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 on page 147

package-name

Syntax package-name *package-name*;

Hierarchy Level [edit virtual-chassis auto-sw-update]

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

Description Specify the software package name or location of the software package to be used by the automatic software update feature for Virtual Chassis or VCF.

Default No package name is specified.

Options ***package-name***—Name of the software package or the URL to the software package to be used.

- 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*

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches*
- Configuring Automatic Software Update on Virtual Chassis Member Switches on page 116
- *Understanding Software Upgrades in a Virtual Chassis Fabric*

preprovisioned

Syntax

```
preprovisioned;
```

Hierarchy Level

[edit virtual-chassis]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 mastership 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.

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 a QFX Series Virtual Chassis on page 86
- Replacing a Member Switch of a Virtual Chassis Configuration on page 103

rcp-count

Syntax rcp-count *number*;

Hierarchy Level [edit chassis (EX Series) nssu]

Release Information Statement introduced in Junos OS Release 14.1X53-D40 for QFX5100 switches.

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 master 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 aborts 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.

Related Documentation

- Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis on page 142
- *Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric*

role

Syntax role (line-card | routing-engine);

Hierarchy Level [edit virtual-chassis member *member-id*]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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" on page 50 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 mastership priority or member ID of member switches through the user interfaces. The mastership priority value is generated by the software, based on the assigned role:

- A member configured as **routing-engine** is assigned the mastership priority 129.
- A member configured as **line-card** is assigned the mastership priority 0.
- A member listed in the preprovisioned configuration without an explicitly specified role is assigned the **line-card** role and mastership priority 0 by default.

The configured role specifications are permanent. If both **routing-engine** members fail, a **line-card** member cannot take over as master 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 mastership if the members functioning as master 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 master 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 master or backup Routing Engine of the Virtual Chassis or VCF configuration. The master manages all members and runs the chassis management processes and control protocols. The backup synchronizes with the master 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 master 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 master, based on the master election algorithm. See "Understanding How the Master in a Virtual Chassis Is Elected" on page 57. 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" on page 50 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.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis on page 81
- *Configuring an EX9200 Virtual Chassis*
- Configuring a QFX Series Virtual Chassis on page 86
- Replacing a Member Switch of a Virtual Chassis Configuration on page 103

serial-number

Syntax serial-number *serial-number*;

Hierarchy Level [edit virtual-chassis member *member-id*]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis on page 81
- *Configuring an EX9200 Virtual Chassis*
- Configuring a QFX Series Virtual Chassis on page 86

serial-number (Virtual Chassis aliases)

Syntax

```
serial-number serial-number {
  alias-name alias-name;
}
```

Hierarchy Level

[edit virtual-chassis aliases]

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series Virtual Chassis and Virtual Chassis Fabric (VCF).

Description

Specify the serial number that will be labeled with an alias in a Virtual Chassis or Virtual Chassis Fabric (VCF). Serial number values are case-sensitive.

The remaining statements are explained separately. See CLI Explorer.

Options

serial-number—Permanent serial number for the member switch in the Virtual Chassis or VCF.

You can retrieve the serial number for any device in your Virtual Chassis or VCF by entering the **show virtual-chassis** command and reviewing the output in the **Serial No** field.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- Configuring a QFX Series Virtual Chassis on page 86
- *Understanding Virtual Chassis Fabric Components*
- Understanding Virtual Chassis Components on page 37

traceoptions (Virtual Chassis)

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]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Option **detail** added in Junos OS Release 9.2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 master 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.

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 on page 126
- Verifying That Virtual Chassis Ports Are Operational on page 127
- Troubleshooting an EX Series Virtual Chassis on page 131
- *Troubleshooting Virtual Chassis Fabric*

upgrade-group

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 (EX Series) nssu]

Release Information

Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.

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; QFX3500, QFX3600, and QFX5100 Virtual Chassis; and Virtual Chassis Fabric (VCF).

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.

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 on page 147

vc-port

Syntax

```
vc-port {
  lag-hash (packet-based | source-port-based);
}
```

Hierarchy Level

[edit virtual-chassis]

Release Information

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Enable hashing of link aggregation group (LAG) network traffic over a dedicated trunk port within a Virtual Chassis.

You select how to direct all LAG traffic through the dedicated trunk port by using the **lag-hash** statement.

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.**

The remaining statement is explained separately. See CLI Explorer.

Default

source-port-based

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

vcp-no-hold-time

Syntax vcp-no-hold-time;

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

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.

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: For the EX4300 Virtual Chassis, 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 62

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.

Related Documentation

- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- Understanding Virtual Chassis Components on page 37

virtual-chassis

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]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 mastership priority of 128, and a default role as master.

A QFX Series device configured in standalone mode is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

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.

Related Documentation

- *Autoprovisioning a Virtual Chassis Fabric*
- *Preprovisioning a Virtual Chassis Fabric*
- *Adding a Device to a Virtual Chassis Fabric*
- Configuring a QFX Series Virtual Chassis on page 86
- Configuring an EX2300, EX3400, or EX4300 Virtual Chassis on page 71
- Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis on page 81
- *Configuring an EX9200 Virtual Chassis*

CHAPTER 6

Operational Commands

- clear virtual-chassis vc-port statistics
- request session member
- request system software nonstop-upgrade
- request virtual-chassis mode
- request virtual-chassis recycle
- request virtual-chassis renumber
- request virtual-chassis vc-port
- request virtual-chassis vc-port diagnostics optics
- show chassis nonstop-upgrade
- show virtual-chassis
- show virtual-chassis active-topology
- show virtual-chassis device-topology
- show virtual-chassis login
- show virtual-chassis mode
- show virtual-chassis protocol adjacency
- show virtual-chassis protocol database
- show virtual-chassis protocol interface
- show virtual-chassis protocol route
- show virtual-chassis protocol statistics
- show virtual-chassis vc-path
- show virtual-chassis vc-port
- show virtual-chassis vc-port diagnostics optics
- show virtual-chassis vc-port statistics

clear virtual-chassis vc-port statistics

Syntax

```
clear virtual-chassis vc-port statistics
<all-members>
<interface-name>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
The options **all-members** and **local** were added in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric.

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

Related Documentation

- show virtual-chassis vc-port statistics on page 284
- show virtual-chassis vc-port on page 264
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*

List of Sample Output

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
```

request session member

Syntax **request session member *member-id***

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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

Related Documentation
- member on page 168
- Understanding Virtual Chassis Components on page 37

request system software nonstop-upgrade

Syntax

```
request system software nonstop-upgrade (package-name | set [package-name
  package-name])
<force-host>
<no-copy>
<no-old-master-upgrade>
<reboot >
<unlink>
```

Release Information

Command introduced in Junos OS Release 10.4 for EX Series switches.
Option **set [*package-name package-name*]** added in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D20 for the QFX Series.
Command introduced in Junos OS Release 15.1X53-D55 for EX3400 switches.

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 on which it is executed, as follows:

- When you execute this command on any of the following Virtual Chassis or VCF configurations, all members are upgraded:
  - EX3300, EX3400, EX4200, EX4300, EX4500, EX4550, or EX4600 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
  - 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 master. The original master is automatically upgraded and rebooted and rejoins the Virtual Chassis or VCF as the backup after the upgrade completes.
- When you execute this command on an EX6200 or EX8200 switch, both the backup and master Routing Engines are upgraded, with the original backup Routing Engine becoming the new master at the end of the upgrade.
  - On an EX6200 switch, the original master Routing Engine is automatically rebooted.
  - On an EX8200 switch, the original master Routing Engine is not automatically rebooted unless you specify the **reboot** option.
- When you execute this command on an EX8200 Virtual Chassis, all master and backup Routing Engines are upgraded in the Virtual Chassis, including the external Routing Engines. The original backup Routing Engines become the new master Routing Engines.

The original master Routing Engines are not automatically rebooted, 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.
- Graceful Routing Engine switchover (GRES) must be enabled.
- Nonstop active routing (NSR) must be enabled.

NOTE: Although nonstop bridging (NSB) does not have to be enabled for you to use this command, we recommend that you enable NSB. Enabling NSB 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 EX Series Switches (CLI Procedure)*.

- The command must be executed from the master Routing Engine on a standalone switch or from the master 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 members (for EX3300, EX3400, EX4200, EX4300, EX4500, EX4550, EX4600, QFX3500 and QFX3600 Virtual Chassis, and mixed Virtual Chassis, and VCF) or on different line cards (for EX6200 and EX8200 switches, and for EX8200 Virtual Chassis).
- For EX3300, EX3400, EX4200, EX4300, EX4500, EX4550, EX4600, QFX3500 and QFX3600 Virtual Chassis, and mixed 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 master and backup must be adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in line-card roles are rebooting.
  - The Virtual Chassis must be pre-configured so that the line-card role has been explicitly assigned to member switches acting in a line-card role. During an NSSU, the Virtual Chassis members must maintain their roles—the master and backup must maintain their Routing Engine roles (although mastership will change), and the remaining switches must maintain their line-card roles.
  - A two-member Virtual Chassis must have **no-split-detection** configured so that the Virtual Chassis does not split when an NSSU upgrades a member.
- For Virtual Chassis Fabric:

- Only two pre-provisioned members in the Routing Engine role are supported. If more than two Routing Engines are configured, a warning is issued and the NSSU process stops.
- The VCF members are connected in a spine and leaf topology. A spine and leaf topology prevents the VCF from splitting during an NSSU. Each leaf device must be connected to both spine devices.
- The VCF must be pre-confured so that the line-card role has been explicitly assigned to member switches acting in a line-card role, and that the Routing Engine role has been explicitly assigned to member switches acting in a Routing Engine role. During an NSSU, the VCF members must maintain their roles—the master and backup must maintain their master and backup roles (although mastership will change), the member switches must remain their Routing Engine roles, and the remaining switches must maintain their line-card roles.
- A two-member VCF must have **no-split-detection** configured so that the VCF does not split when an NSSU upgrades a member.

Options *package-name*—Location from which the software package or bundle is to be installed. For example:

- **/var/tmp/*package-name***—For a software package or bundle that is being installed from a local directory on the switch.
- ***protocol*://*hostname*/*pathname*/*package-name***—For a software package or bundle that is to be 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. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
  - **http**—Hypertext Transfer Protocol.
    Use **http://*hostname*/*pathname*/*package-name***. To specify authentication credentials, use **http://<*username*>:<*password*>@*hostname*/*pathname*/*package-name***. If a password is required and you omit it, you are prompted for 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 EX4200 and the EX4500 installation packages. These packages must be for the same Junos OS release. See the description of the ***package-name*** option for information about how to specify the location of the installation packages.

force-host—(Optional) Force the addition of host software package or bundle (ignore warnings) on a QFX5100 device.

no-copy—(Optional) Install a software package or bundle, but do not save copies of package or bundle files.

no-old-master-upgrade—(Optional) (EX8200 switches only) Upgrade the backup Routing Engine only. After the upgrade completes, the original master Routing Engine becomes the backup Routing Engine and continues running the previous software version.

reboot—(Optional) (EX8200 switches and EX8200 Virtual Chassis only) When the **reboot** option is included, the original master (new backup) Routing Engines are automatically rebooted after being upgraded to the new software. When the **reboot** option is not included, you must manually reboot the original master (new backup) Routing Engines 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 perform the manual reboot of the backup Routing Engines.

unlink—(Optional) Remove the software package after a successful upgrade is completed.

Required Privilege Level

maintenance

Related Documentation

- show chassis nonstop-upgrade on page 218
- *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 on page 150
- *Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*

List of Sample Output

request system software nonstop-upgrade (EX4200 Virtual Chassis) on page 203
request system software nonstop-upgrade (EX6200 Switch) on page 204
request system software nonstop-upgrade reboot (EX8200 Switch) on page 205

request system software nonstop-upgrade no-old-master-upgrade (EX8200 Switch) on page 206
request system software nonstop-upgrade reboot (EX8200 Virtual Chassis) on page 207

Output Fields When you enter this command, you are provided feedback on the status of your request.

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
```

request virtual-chassis mode

Syntax

```
request virtual-chassis mode
  fabric
  mixed <ieee-clause-82>
  <disable>
  <reboot>
  <all-members>
  <local>
  <member member-id>
```

Release Information

Command introduced in Junos OS Release 11.1 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.
fabric keyword introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices in a Virtual Chassis Fabric (VCF).
Command introduced in Junos OS Release 13.2X51-D20 for 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.

Description

Configure the mode for a device or multiple devices in a Virtual Chassis or a VCF. The Virtual Chassis mode setting is maintained through reboots even though it is set in operational mode.

Fabric Mode

A device is configured in fabric mode to participate as a member device in a VCF.

Mixed Mode

A device must be configured in mixed mode when it is participating in a Virtual Chassis or a VCF with devices that have differences in how they interoperate. See "Understanding Mixed EX Series and QFX Series Virtual Chassis" on page 50 or *Understanding Mixed Virtual Chassis Fabric* for details on which devices can be interconnected to form a mixed Virtual Chassis or VCF.

NOTE: You do not need to configure mixed mode if the only devices in your Virtual Chassis are the following combinations of switches, which form a non-mixed Virtual Chassis or VCF because the devices can run the same software images:

- Only EX4500 and EX4550 switches in an EX Series 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.

You must include the **ieee-clause-82** option when setting **mixed** mode on EX4300 switches that are not multigigabit models for those switches to participate in an EX4300 Virtual Chassis with EX4300 multigigabit model switches (EX4300-48MP). Due to platform differences, this mode is required for other EX4300 switches to properly communicate with EX4300 multigigabit switches.

Do not enable the **request virtual-chassis mode mixed** command for a standalone device or for a member switch that is intended to remain in a non-mixed Virtual Chassis or VCF. Enabling this command reduces the maximum scaling numbers for some features on the switch, Virtual Chassis, or VCF.

BEST PRACTICE: For a mixed Virtual Chassis, to avoid potential traffic disruptions and configuration issues, we recommend configuring mixed mode on your device before cabling it into your Virtual Chassis, and rebooting the device to complete this configuration procedure before interconnecting it into the Virtual Chassis. Similarly, for a VCF, to avoid potential traffic disruptions and configuration issues, we recommend configuring the fabric and, if applicable, the mixed mode settings on your device and rebooting it before cabling it into the VCF. You can change the fabric and mixed mode settings manually after a device has been added to a Virtual Chassis or VCF.

If you set only some of the devices in a mixed Virtual Chassis or VCF to mixed mode using this command, 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 set the Virtual Chassis mode to mixed for all devices in the Virtual Chassis or VCF. You then need to reboot the devices that have been set into mixed mode to complete the procedure. The Virtual Chassis or VCF forms after the devices have rebooted.

NOTE: When you do not use this command to set any of the switches in a mixed EX4200 and EX4500 Virtual Chassis to mixed mode, a mixed EX4200 and EX4500 Virtual Chassis forms with one of the switches assuming the master role if the switches are running Junos OS Release 11.4 or later. All other switches in the mixed EX4200 and EX4500 Virtual Chassis are placed into the linecard role. If you experience this behavior, enter the **request virtual-chassis mode mixed all-members** command to set the Virtual Chassis mode to mixed for all switches in the Virtual Chassis. You will then need to reboot the switches to complete the procedure. The Virtual Chassis will form after all of the switches have rebooted.

Options

No options (default)—Set the Virtual Chassis mode for all members of the Virtual Chassis or VCF.

all-members—(Optional) Set the Virtual Chassis mode for all members of the Virtual Chassis or VCF.

disable—Disable the Virtual Chassis fabric or mixed mode setting if it was previously enabled.

fabric—Set the device into fabric mode so that the device can participate 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. The **ieee-clause-82** mixed-mode option is required when mixing EX4300 switches that are not multigigabit models with EX4300 multigigabit (EX4300-48MP) model switches in an EX4300 Virtual Chassis.

reboot—After applying the mode change specified by the other options, reboot the device automatically.

Required Privilege Level system-control

Related Documentation

- Understanding Mixed EX Series and QFX Series Virtual Chassis on page 50
- *Understanding Mixed Virtual Chassis Fabric*
- *Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)*
- *Verifying the Virtual Chassis Fabric Mode Settings*
- Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 126

List of Sample Output

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
```

request virtual-chassis recycle

Syntax **request virtual-chassis recycle member-id *member-id***

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

Description Make a previously used member ID available for reassignment.

When you remove a member switch from the Virtual Chassis configuration, the master 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 master 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

Related Documentation
- request virtual-chassis renumber on page 214
- Replacing a Member Switch of a Virtual Chassis Configuration on page 103

List of Sample Output
request virtual-chassis recycle member-id 3 on page 213
request virtual-chassis recycle member-id 1 on page 213

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
```

request virtual-chassis renumber

Syntax request virtual-chassis renumber member-id *old-member-id* new-member-id *new-member-id*

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

Description Renumber a member of a Virtual Chassis configuration.

NOTE: You must run this command from the Virtual Chassis member in the master 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

Related Documentation

- request virtual-chassis recycle on page 213
- Replacing a Member Switch of a Virtual Chassis Configuration on page 103

List of Sample Output

request virtual-chassis renumber member-id 5 new-member-id 4 on page 214
request virtual-chassis renumber member-id 1 new-member-id 0 on page 214

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
```

request virtual-chassis vc-port

Syntax

```
request virtual-chassis vc-port [set | delete]
<fpc-slot fpc-slot>
pic-slot pic-slot
port port-number
<member member-id>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **fpc-slot** introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

Description

Enable or disable an optical port as a Virtual Chassis port (VCP).

If you omit **member *member-id***, this command defaults to enabling or disabling the uplink VCP or SFP network port configured as a VCP on the switch where the command is issued.

You might experience a temporary traffic disruption immediately after creating or deleting a user-configured VCP in an EX8200 Virtual Chassis.

NOTE: Some EX Series switches have ports that are configured as VCPs by default. You do not need to explicitly configure those ports as VCPs to use them to interconnect the switch into a Virtual Chassis. See "Understanding Virtual Chassis Components" on page 37 for details.

Options

set—Set a network port as a VCP to convert a network port into a VCP.

delete—Delete the VCP setting on a port to convert a VCP into a 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 is to be enabled or disabled as a VCP.

member *member-id*—(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis or VCF.

Required Privilege Level

system-control

Related Documentation

- *request virtual-chassis vc-port (Dedicated VCP)*
- show virtual-chassis vc-port on page 264
- show virtual-chassis vc-port statistics on page 284

- clear virtual-chassis vc-port statistics on page 196
- *Virtual Chassis Port (VCP) Interface Names in an EX8200 Virtual Chassis*
- Understanding Virtual Chassis Components on page 37

List of Sample Output

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.

request virtual-chassis vc-port diagnostics optics

Syntax **request virtual-chassis vc-port diagnostics optics**

Release Information Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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, 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.

Required Privilege Level system-control

Related Documentation
- show virtual-chassis vc-port diagnostics optics on page 269

Sample Output

request virtual-chassis vc-port diagnostics optics

```
user@switch> request virtual-chassis vc-port diagnostics optics
fpc0:
--------------------------------------------------------------------------
vc-port Diagnostics Optics Done
```

show chassis nonstop-upgrade

Syntax

show chassis nonstop-upgrade

Release Information

Command introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

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). This command must be issued on the master Routing Engine.

Required Privilege Level

view

Related Documentation

- request system software nonstop-upgrade on page 199
- *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 on page 150
- *Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*
- *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*

List of Sample Output

show chassis nonstop-upgrade (EX8200 Switch) on page 219
show chassis nonstop-upgrade (EX8200 Virtual Chassis) on page 219
show chassis nonstop-upgrade (Virtual Chassis Fabric) on page 219

Output Fields

Table 9 on page 218 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
```

show virtual-chassis

Syntax **show virtual-chassis <status>**

Release Information Command introduced in Junos OS Release 9.2 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
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.

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

Related Documentation

- show virtual-chassis active-topology on page 226
- show virtual-chassis protocol adjacency on page 242
- show virtual-chassis vc-path on page 260
- Understanding Mixed EX Series and QFX Series Virtual Chassis on page 50
- *Understanding Mixed Virtual Chassis Fabric*
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*

List of Sample Output

show virtual-chassis (EX2300 multigigabit model—EX2300-24MP and EX2300-48MP—Virtual Chassis) on page 222
show virtual-chassis (EX2300 and EX2300 multigigabit model (EX2300-24MP, EX2300-48MP) Virtual Chassis) on page 223
show virtual-chassis (EX4200 Virtual Chassis) on page 223
show virtual-chassis (Mixed EX4300 multigigabit model—EX4300-48MP—Virtual Chassis) on page 223
show virtual-chassis (EX8200 Virtual Chassis) on page 224
show virtual-chassis (QFX5110 Virtual Chassis) on page 224
show virtual-chassis (QFX5200 Virtual Chassis) on page 225
show virtual-chassis (QFX5100 Virtual Chassis Fabric) on page 225

Output Fields Table 10 on page 221 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.
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.
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.

Table 10: show virtual-chassis Output Fields (continued)

Field Name	Field Description
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**	Mastership priority value of the member device.
Role	Role of the member device: master, backup, or linecard. An asterisk (*) following the Role denotes the member device on which the **show virtual-chassis <status>** command was issued.
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: fabric (F) or Virtual Chassis (V).
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   XY1222340011 ex2300-24mp  128   Backup     N     VC   3  vcp-255/1/0
                                                                    1  vcp-255/1/2
1 (FPC 1)  Prsnt   XY1222340021 ex2300-48mp  128   Master*    N     VC   2  vcp-255/1/4
                                                                    0  vcp-255/1/1
2 (FPC 2)  Prsnt   XY1222340031 ex2300-24mp  128   Linecard   N     VC   3  vcp-255/1/0
                                                                    1  vcp-255/1/1
3 (FPC 3)  Prsnt   XY1222340041 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    XP0123450001 ex2300-24mp    129   Backup       N  VC   1  vcp-255/1/3
                                                                          3  vcp-255/1/0
1 (FPC 1)  Prsnt    JJ0123450001 ex2300-24p     129   Master*      N  VC   2  vcp-255/1/0
                                                                          0  vcp-255/1/1
2 (FPC 2)  Prsnt    JJ0123450002 ex2300-24p       0   Linecard     N  VC   3  vcp-255/1/2
                                                                          1  vcp-255/1/3
3 (FPC 3)  Prsnt    XP0123450002 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   AK0207360276 ex4200-24t     249   Master*    N   8   vcp-0
                                                                 1   vcp-1
1 (FPC 1)  Prsnt   AK0207360281 ex4200-24t     248   Backup     N   0   vcp-0
                                                                 2   vcp-1
2 (FPC 2)  Prsnt   AJ0207391130 ex4200-48p     247   Linecard   N   1   vcp-0
                                                                 3   vcp-1
3 (FPC 3)  Prsnt   AK0207360280 ex4200-24t     246   Linecard   N   2   vcp-0
                                                                 4   vcp-1
4 (FPC 4)  Prsnt   AJ0207391113 ex4200-48p     245   Linecard   N   3   vcp-0
                                                                 5   vcp-1
5 (FPC 5)  Prsnt   BP0207452204 ex4200-48t     244   Linecard   N   4   vcp-0
                                                                 6   vcp-1
6 (FPC 6)  Prsnt   BP0207452222 ex4200-48t     243   Linecard   N   5   vcp-0
                                                                 7   vcp-1
7 (FPC 7)  Prsnt   BR0207432028 ex4200-24f     242   Linecard   N   6   vcp-0
                                                                 8   vcp-1
8 (FPC 8)  Prsnt   BR0207431996 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    XR3617480028 ex4300-48mp    129   Master*      Y  VC   1  vcp-255/1/0
```

```
                                                                    3  vcp-255/1/3
1 (FPC 1)  Prsnt    XR3617480029 ex4300-48mp    129   Backup       Y  VC   0  vcp-255/1/1
                                                                    2  vcp-255/1/3
2 (FPC 2)  Prsnt    XR3617480001 ex4300-48mp      0   Linecard     Y  VC   3  vcp-255/1/1
                                                                    1  vcp-255/1/3
3 (FPC 3)  Prsnt    PE3715471074 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   BA0908380001 ex8216          0  Linecard   8  vcp-0/0
                                                                   8  vcp-0/1
                                                                   1  vcp-4/0/4
1 (FPC 16-31)   Prsnt   BT0909411634 ex8208          0  Linecard   8  vcp-0/0
                                                                   0  vcp-3/0/4
8 (FPC 128-143) Prsnt   062009000021 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   062009000022 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
Member ID  Status   Serial No    Model         prio  Role      Mode  Mode ID Interface
0 (FPC 0)  Prsnt    WT3116360038 qfx5110-32q   129   Backup       Y  VC   9  vcp-255/0/12
                                                                    1  vcp-255/0/16
1 (FPC 1)  Prsnt    WT3116360041 qfx5110-32q   129   Master*      Y  VC   2  vcp-255/0/16
                                                                    0  vcp-255/0/17
2 (FPC 2)  Prsnt    WT3116360051 qfx5110-32q     0   Linecard     Y  VC   3  vcp-255/0/11
                                                                    1  vcp-255/0/14
3 (FPC 3)  Prsnt    WT3116360055 qfx5110-32q     0   Linecard     Y  VC   4  vcp-255/0/14
                                                                    2  vcp-255/0/31
4 (FPC 4)  Prsnt    WT3116360056 qfx5110-32q     0   Linecard     Y  VC   5  vcp-255/0/10
                                                                    3  vcp-255/0/5
5 (FPC 5)  Prsnt    TB3716340058 qfx5100e-24q-2p 0   Linecard     Y  VC   6  vcp-255/0/12
                                                                    4  vcp-255/0/20
6 (FPC 6)  Prsnt    WS3715500018 qfx5110-48s-4c  0   Linecard     Y  VC   7  vcp-255/0/10
                                                                    5  vcp-255/0/49
7 (FPC 7)  Prsnt    TA3714110007 qfx5100e-48s-6q 0   Linecard     Y  VC   6  vcp-255/0/18
```

```
                                                                                  8  vcp-255/0/31
8 (FPC 8)   Prsnt    WS3715500022 qfx5110-48s-4c   0   Linecard     Y  VC    7  vcp-255/0/21
                                                                                  9  vcp-255/0/49
9 (FPC 9)   Prsnt    WT3116360061 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
                                                                  Mixed Route Neighbor List
Member ID  Status   Serial No    Model           prio  Role      Mode  Mode  ID Interface
0 (FPC 0)  Prsnt    XY0123456789 qfx5200-32c-r   128   Master*   N     VC    1  vcp-255/0/19
                                                                                2  vcp-255/0/26
1 (FPC 1)  Prsnt    XY0123456780 qfx5200-32c-r   128   Linecard  N     VC    0  vcp-255/0/20
                                                                                2  vcp-255/0/16
2 (FPC 2)  Prsnt    YZ0123456789 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
                                                Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model          prio  Role      Mode Mode  ID Interface
0 (FPC 0)  Prsnt    AB3112430001 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    AB3112230001 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    AB3112460011 qfx5100-48s      0   Linecard     N  F     1  vcp-255/1/0
                                                                            0  vcp-255/1/1
3 (FPC 3)  Prsnt    AB3112460011 qfx5100-48s      0   Linecard     N  F     1  vcp-255/1/0
                                                                            0  vcp-255/1/1
4 (FPC 4)  Prsnt    AB3112430011 qfx5100-48s      0   Linecard     N  F     1  vcp-255/1/0
                                                                            0  vcp-255/1/1
```

show virtual-chassis active-topology

Syntax

```
show virtual-chassis active-topology
<all-members>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 the command is issued.

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 on which this command is entered.

member *member-id*—(Optional) Display the active topology of the specified member of the Virtual Chassis or VCF.

Required Privilege Level

view

Related Documentation

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

List of Sample Output

show virtual-chassis active-topology (EX4200 Virtual Chassis) on page 227
show virtual-chassis active-topology (EX8200 Virtual Chassis) on page 227
show virtual-chassis active-topology (Virtual Chassis Fabric) on page 228

Output Fields

Table 11 on page 226 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)
```

show virtual-chassis device-topology

Syntax

```
show virtual-chassis device-topology
<all-members>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 on which this command is entered.

member *member-id*—(Optional) Display the device topology for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Related Documentation

- Understanding Virtual Chassis Port Link Aggregation on page 61
- *Understanding EX8200 Virtual Chassis Topologies*

Output Fields

Table 12 on page 231 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.

Table 12: show virtual-chassis device-topology Output Fields (continued)

Field Name	Field Description
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
```

show virtual-chassis login

Syntax

```
show virtual-chassis login
```

Release Information

Command introduced in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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

Related Documentation

- request session member on page 198
- Understanding Global Management of a Virtual Chassis on page 58

List of Sample Output

show virtual-chassis login (Direct Login to the Master Console Port) on page 238
show virtual-chassis login (Backup Console Session Redirected to the Master Console Port) on page 238

Sample Output

show virtual-chassis login (Direct Login to the Master 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 Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host backup
```

show virtual-chassis mode

Syntax

```
show virtual-chassis mode
<all-members>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 11.1 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
Current mode and **Future mode after reboot** fields introduced in Junos OS Release 13.2X51-D20.

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 on which the command is entered.

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 on which the command is entered.

member *member-id*—(Optional) Display the Virtual Chassis or VCF mixed mode status for the specified member device..

Required Privilege Level

view

Related Documentation

- request virtual-chassis mode on page 209
- *Verifying the Virtual Chassis Fabric Mode Settings*
- *Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)*

List of Sample Output

show virtual-chassis mode (EX4200) on page 240
show virtual-chassis mode (QFX5100) on page 240

Output Fields

Table 13 on page 239 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	Specifies the current mixed and fabric mode settings running on the member device or devices. A device reboot is required to change the fabric or mixed 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.
Future mode after reboot	Specifies the mixed and fabric mode settings running on the member device or devices. A device reboot is required to change the fabric or mixed 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.

Sample Output

show virtual-chassis mode (EX4200)

```
user@switch>show virtual-chassis mode

fpc0:
--------------------------------------------------------------------------
Mixed Mode: Disabled
```

Sample Output

show virtual-chassis mode (QFX5100)

```
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
```

show virtual-chassis protocol adjacency

Syntax

```
show virtual-chassis protocol adjacency
<brief | detail | extensive>
<all-members>
<local>
<member member-id>
<system-id>
```

Release Information

Command introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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

Related Documentation

- Understanding Virtual Chassis Port Link Aggregation on page 61
- *Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis*

List of Sample Output

show virtual-chassis protocol adjacency on page 243
show virtual-chassis protocol adjacency detail on page 244

Output Fields

Table 14 on page 243 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
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
```

show virtual-chassis protocol database

Syntax

```
show virtual-chassis protocol database
<brief | detail | extensive>
<all-members>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 on which this command is entered.

member *member-id*—(Optional) Display VCCP database statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Related Documentation

- *Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis*
- Understanding Virtual Chassis Components on page 37

List of Sample Output

show virtual-chassis protocol database on page 247
show virtual-chassis protocol database detail on page 248

Output Fields

Table 15 on page 247 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
```

show virtual-chassis protocol interface

Syntax

```
show virtual-chassis protocol interface
<brief | detail>
<all-members>
<interface-name>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 on which this command is entered.

member *member-id*—(Optional) Display VCCP interface statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Related Documentation

- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- *Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis*
- *Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis*

List of Sample Output

show virtual-chassis protocol interface on page 251

Output Fields

Table 16 on page 251 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
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
```

show virtual-chassis protocol route

Syntax

```
show virtual-chassis protocol route
<all-members>
<destination-id>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 this command is entered.

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

Related Documentation

- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- *Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis*

List of Sample Output

show virtual-chassis protocol route on page 254

Output Fields

Table 17 on page 253 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.

Table 17: show virtual-chassis protocol route Output Fields (continued)

Field Name	Field Description
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.
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
```

show virtual-chassis protocol statistics

Syntax

```
show virtual-chassis protocol statistics
<all-members>
<interface-name>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 on which this command is entered.

member *member-id*—(Optional) Display VCCP statistics for the specified member of the Virtual Chassis or VCF.

Required Privilege Level

clear

Related Documentation

- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- *Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis*

List of Sample Output

show virtual-chassis protocol statistics on page 257

Output Fields

Table 18 on page 256 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.

Table 18: show virtual-chassis protocol statistics Output Fields (continued)

Field Name	Field Description
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.
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
```

show virtual-chassis vc-path

Syntax

show virtual-chassis vc-path source-interface ***interface-name*** **destination-interface** ***interface-name***

Release Information

Command introduced in Junos OS Release 9.6 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

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

Related Documentation

- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*
- Understanding EX Series Virtual Chassis on page 24
- Understanding QFX Series Virtual Chassis on page 32
- *EX8200 Virtual Chassis Overview*

List of Sample Output

show virtual-chassis vc-path source-interface destination-interface (Virtual Chassis) on page 261
show virtual-chassis vc-path source-interface destination-interface (Virtual Chassis Fabric) on page 262

Output Fields

Table 19 on page 261 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 261 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.
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  VX1234560001 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  VX1234560002 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  VX1234560003 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
4 (FPC 4)  Prsnt  VX1234560004 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
5 (FPC 5)  Prsnt  VX1234560005 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
6 (FPC 6)  Prsnt  VX1234560006 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
7 (FPC 7)  Prsnt  VX1234560007 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
8 (FPC 8)  Prsnt  VX1234560008 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
9 (FPC 9)  Prsnt  VX1234560009 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
10 (FPC 10)Prsnt  VX1234560010 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
11 (FPC 11)Prsnt  VX1234560011 qfx5100-24q-aa   0   Linecard    N  F    1  vcp-255/0/1
                                                                      2  vcp-255/0/2
12 (FPC 12)Prsnt  VX1234560012 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
```

show virtual-chassis vc-port

Syntax

```
show virtual-chassis vc-port
<all-members>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

Description

Display the status of the Virtual Chassis ports (VCPs), including both dedicated VCPs and uplink ports configured as VCPs, if present.

Options

none—Display the operational status of all VCPs of the member switch where the command is issued.

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 on which this command is entered.

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

Related Documentation

- show virtual-chassis vc-port statistics on page 284
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*

List of Sample Output

show virtual-chassis vc-port (EX4200 Virtual Chassis) on page 266
show virtual-chassis vc-port (EX8200 Virtual Chassis) on page 266
show virtual-chassis vc-port all-members on page 267

Output Fields

Table 20 on page 264 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
fpc*number*	The FPC number is the same as the member ID.

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 instance, a display of **vcp-255/1/0** indicates that the dedicated VCP named **vcp-1/0** is part of a LAG bundle. A display of **vcp-255/1/0** indicates that an uplink port that was previously named **xe-0/1/0** is now part of a VCP LAG bundle.
Type	Type of VCP: • **Dedicated**—The rear panel VCP on an EX4200, 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. • **Auto-Configured**—Optical port autoconfigured as a VCP. See "Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port" on page 110 or "Configuring a QFX Series Virtual Chassis" on page 86 for information about configuring VCPs.
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 diagnostics optics

Syntax

```
show virtual-chassis vc-port diagnostics optics
<all-members>
<interface-name>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 12.2 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 on which this command is entered.

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

Related Documentation

- show virtual-chassis vc-port on page 264
- *Installing a Transceiver*
- *Removing a Transceiver*
- *Junos OS Ethernet Interfaces Configuration Guide*

List of Sample Output

Output Fields

Table 21 on page 270 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.
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*.

Table 21: show virtual-chassis vc-port diagnostics optics Output Fields (continued)

Field Name	Field Description
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*.
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.
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.

Table 21: show virtual-chassis vc-port diagnostics optics Output Fields (continued)

Field Name	Field Description
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.
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
```

show virtual-chassis vc-port statistics

Syntax

```
show virtual-chassis vc-port statistics
<all-members>
<brief | detail | extensive >
<interface-name>
<local>
<member member-id>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
The options **all-members**, **brief**, **detail**, **extensive**, and **local** were added in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

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 on which this command is entered.

member *member-id*—(Optional) Display traffic statistics for VCPs on the specified member of a Virtual Chassis or VCF.

Required Privilege Level

view

Related Documentation

- clear virtual-chassis vc-port statistics on page 196
- show virtual-chassis vc-port on page 264
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*
- *Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis*

List of Sample Output

show virtual-chassis vc-port statistics member 0 on page 290

Output Fields Table 21 on page 270 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 master 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**
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**

Table 22: show virtual-chassis vc-port statistics Output Fields (continued)

Field Name	Field Description	Level of Output
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**
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}
```