# Junos® OS for EX Series Ethernet Switches

## High Availability for EX Series Switches

Release
12.3

Published: 2013-01-03

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net



*Junos® OS for EX Series Ethernet Switches High Availability for EX Series Switches*
Release 12.3


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 http://www.juniper.net/support/eula.html. 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 http://www.juniper.net/techpubs/.

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 http://www.juniper.net/books.

## Supported Platforms

For the features described in this document, the following platforms are supported:

- EX Series

## 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 the CLI User Guide.

## Documentation Conventions

Table 1 on page xiii 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. |

Table 2 on page xiii 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:<br>user@host> **configure** |
| `Fixed-width text like this` | Represents output that appears on the terminal screen. | `user@host>` **show chassis alarms**<br>`No alarms currently active` |

Table 2: Text and Syntax Conventions *(continued)*

| Convention | Description | Examples |
|---|---|---|
| *Italic text like this* | • Introduces or emphasizes important new terms.<br>• Identifies book names.<br>• Identifies RFC and Internet draft titles. | • A policy *term* is a named structure that defines match conditions and actions.<br>• *Junos OS System Basics Configuration Guide*<br>• 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:<br>[edit]<br>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.<br>• The console port is labeled **CONSOLE**. |
| < > (angle brackets) | Enclose 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**<br>**(*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) | Enclose a variable for which you can substitute one or more values. | **community name members [** ***community-ids*** **]** |
| Indention and braces ( { } ) | Identify a level in the configuration hierarchy. | [edit]<br>routing-options {<br>  static {<br>  route default {<br>  nexthop *address*;<br>  retain;<br>  }<br>  }<br>} |
| ; (semicolon) | Identifies a leaf statement at a configuration hierarchy level. | |
| J-Web GUI Conventions | | |
| **Bold text like this** | Represents J-Web graphical user interface (GUI) items you click or select. | • In the Logical Interfaces box, select **All Interfaces**.<br>• To cancel the configuration, click **Cancel**. |
| **>** (bold right angle bracket) | Separates levels in a hierarchy of J-Web selections. | In the configuration editor hierarchy, select **Protocols>Ospf**. |

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release 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 JNASC 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 http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://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: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/

- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

## Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- 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 http://www.juniper.net/support/requesting-support.html.

PART 1

# Overview

 

CHAPTER 1

# High Availability Features Overview



## High Availability Features for EX Series Switches Overview

*High availability* refers to the hardware and software components that provide redundancy and reliability for packet-based communications. This topic covers the following high availability features of Juniper Networks EX Series Ethernet Switches:



### VRRP

You can configure Virtual Router Redundancy Protocol (VRRP) for IP and IPv6 on Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and logical interfaces on the switches. When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts.

## Graceful Protocol Restart

With standard implementations of routing protocols, any service interruption requires an affected switch to recalculate adjacencies with neighboring switches, restore routing table entries, and update other protocol-specific information. An unprotected restart of a switch can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. Graceful protocol restart allows a restarting switch and its neighbors to continue forwarding packets without disrupting network performance. Because neighboring switches assist in the restart (these neighbors are called helper switches), the restarting switch can quickly resume full operation without recalculating algorithms from scratch.

On the switches, graceful protocol restart can be applied to aggregate and static routes and for routing protocols (BGP, IS-IS, OSPF, and RIP).

Graceful protocol restart works similarly for the different routing protocols. The main benefits of graceful protocol restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful protocol restart thus allows a switch to pass through intermediate convergence states that are hidden from the rest of the network. Most graceful restart implementations define two types of switches—the restarting switch and the helper switch. The restarting switch requires rapid restoration of forwarding state information so that it can resume the forwarding of network traffic. The helper switch assists the restarting switch in this process. Individual graceful restart configuration statements typically apply to either the restarting switch or the helper switch.

## Redundant Routing Engines

Redundant Routing Engines are two Routing Engines that are installed in a switch or a Virtual Chassis. When a switch has two Routing Engines, one functions as the master, while the other stands by as a backup should the master Routing Engine fail. When a Virtual Chassis has two Routing Engines, the switch in the master role functions as the master Routing Engine and the switch in the backup role functions as the backup Routing Engine. Redundant Routing Engines are supported on Juniper Networks EX6200 Ethernet Switches, Juniper Networks EX8200 Ethernet Switches, and on all EX Series Virtual Chassis configurations.

The master Routing Engine receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding Engine components of the switch, and has full control over the control plane of the switch.

The backup Routing Engine stays in sync with the master Routing Engine in terms of protocol states, forwarding tables, and so forth. If the master becomes unavailable, the backup Routing Engine takes over the functions that the master Routing Engine performs.

Network reconvergence takes place more quickly on switches and on Virtual Chassis with redundant Routing Engines than on switches and on Virtual Chassis with a single Routing Engine.

## Virtual Chassis

A Virtual Chassis is multiple switches connected together that operate as a single network entity. The advantages of connecting multiple switches into a Virtual Chassis include better-managed bandwidth at a network layer, simplified configuration and maintenance because multiple devices can be managed as a single device, and a simplified Layer 2 network topology that minimizes or eliminates the need for loop prevention protocols such as Spanning Tree Protocol (STP). A Virtual Chassis improves high availability by introducing a variety of failover mechanisms; if a member switch, a line card, or an interface fails on a switch that is a member of a Virtual Chassis fails, for instance, traffic to that switch, line card, or interface can be rerouted within the Virtual Chassis.

Juniper Networks EX2200 Ethernet Switches, EX3300 Ethernet Switches, Juniper Networks EX4200 Ethernet Switches, Juniper Networks EX4500 Ethernet Switches, or EX8200 switches can form a Virtual Chassis. EX4200, EX4500, and EX4550 switches can be interconnected together to form a mixed Virtual Chassis.

## Graceful Routing Engine Switchover

You can configure graceful Routing Engine switchover (GRES) on a switch with redundant Routing Engines or on a Virtual Chassis, allowing control to switch from the master Routing Engine to the backup Routing Engine with minimal interruption to network communications. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and forwarding state. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.

When the backup Routing Engine assumes mastership in a redundant failover configuration (that is, when graceful Routing Engine switchover is not enabled), the Packet Forwarding Engines initialize their state to the boot state before they connect to the new master Routing Engine. In contrast, in a graceful switchover configuration, the Packet Forwarding Engines do not reinitialize their state, but resynchronize their state to that of the new master Routing Engine. The interruption to traffic is minimal.

## Link Aggregation

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle. 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 should fail, the system automatically load-balances traffic across all remaining links. In a Virtual Chassis, LAGs can be used to load-balance network traffic between member switches.

The number of Ethernet interfaces you can include in a LAG and the number of LAGs you can configure on a switch depend on the switch model. For information on LAGs, see Understanding Aggregated Ethernet Interfaces and LACP.

## Nonstop Active Routing

Nonstop active routing (NSR) provides high availability in a switch with redundant Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbor routing devices, which do not detect that a change has occurred.

To use nonstop active routing, you must also configure graceful Routing Engine switchover.

## Nonstop Software Upgrade

Nonstop software upgrade (NSSU) allows you to upgrade the software on a switch with dual Routing Engines or on a Virtual Chassis in an automated manner with minimal traffic disruption. NSSU takes advantage of graceful Routing Engine switchover and nonstop active routing to enable upgrading the Junos OS version with no disruption to the control plane. In addition, NSSU minimizes traffic disruption by:

- Upgrading line cards one at a time in an EX8200 switch or EX8200 Virtual Chassis, permitting traffic to continue to flow through the line cards that are not being upgraded.
- Upgrading member switches one at a time in an EX4200, EX4500, or mixed EX4200 and EX4500 Virtual Chassis, permitting traffic to continue to flow through the members that are not being upgraded.

By configuring LAGs such that the member links reside on different line cards or Virtual Chassis members, you can achieve minimal traffic disruption when performing an NSSU.

## Redundant Power System

Most Juniper Networks Ethernet Switches have a built-in capability for redundant power supplies—therefore if one power supply fails on those switches, the other power supply takes over. However, EX2200 switches and EX3300 switches have only one internal fixed power supply. If an EX2200 switch or EX3300 switch is deployed in a critical situation, we recommend that you connect a Redundant Power System (RPS) to that switch to supply backup power if the internal power supply fails. RPS is not a primary power supply—it only provides backup power to switches when the single dedicated power supply fails. An RPS operates in parallel with the single dedicated power supplies of the switches connected to it and provides all connected switches enough power to support either power over Ethernet (PoE) or non-PoE devices. For more information on RPS, see "EX Series Redundant Power System Hardware Overview" on page 29.

**Related Documentation**

- For more information on high availability features, see the *Junos OS High Availability Configuration Guide*.
- EX Series Virtual Chassis Overview
- EX8200 Virtual Chassis Overview
- Understanding VRRP on EX Series Switches on page 35
- Understanding Aggregated Ethernet Interfaces and LACP
- Understanding Nonstop Active Routing on EX Series Switches on page 9

- Understanding Nonstop Software Upgrade on EX Series Switches on page 13
- EX Series Redundant Power System (RPS) Documentation

CHAPTER 2

# Nonstop Active Routing Overview



## Understanding Nonstop Active Routing on EX Series Switches

You can configure nonstop active routing (NSR) on an EX Series switch with redundant Routing Engines to enable the transparent switchover of the Routing Engines in the event that one of the Routing Engines goes down.

Nonstop active routing provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbor routing devices, which do not detect that a change has occurred.

Enable nonstop active routing when neighbor routing devices are not configured to support graceful restart of protocols or when you want to ensure graceful restart of protocols for which graceful restart is not supported (such as PIM).

You do not need to start the two Routing Engines simultaneously to synchronize them for nonstop active routing. If both Routing Engines are not present or not up when you issue a **commit synchronize** statement, the candidate configuration is committed in the master Routing Engine and when the backup Routing Engine is inserted or comes online, its configuration is automatically synchronized with that of the master.

Nonstop active routing uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop active routing also saves routing protocol information by running the routing protocol process (**rpd**) on the backup Routing Engine. By saving this additional information, nonstop active routing does not rely on other routing devices to assist in restoring routing protocol information.

NOTE: After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics** (*interface-name* | **all**) command to reset the cumulative values for local statistics on the new master Routing Engine.

If you suspect a problem with the synchronization of Routing Engines when nonstop active routing is enabled, you can gather troubleshooting information using trace options.

For example, if certain protocols lose connectivity with neighbors after a graceful Routing Engine switchover with NSR enabled, you can use trace options to help isolate the problem. See "Tracing Nonstop Active Routing Synchronization Events" on page 159.

NOTE: Graceful restart and nonstop active routing are mutually exclusive. You will receive an error message upon commit if both are configured.

NOTE: Nonstop active routing provides a transparent switchover mechanism only for Layer 3 protocol sessions. Nonstop bridging (NSB) provides a similar mechanism for Layer 2 protocol sessions. See "Understanding Nonstop Bridging on EX Series Switches" on page 11.

**Related Documentation**

- Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure) on page 47
- Example: Configuring Nonstop Active Routing on EX Series Switches on page 41

CHAPTER 3

# Nonstop Bridging Overview



## Understanding Nonstop Bridging on EX Series Switches

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on a Juniper Networks EX Series Ethernet Switch with redundant Routing Engines.

Nonstop bridging operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because all session information is already synchronized to the backup Routing Engine. Traffic disruption for the NSB-supported Layer 2 protocol is minimal or nonexistent as a result of the switchover. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the NSB-supported Layer 2 protocol sessions on the switch.

For a list of the EX Series switches and Layer 2 protocols that support nonstop bridging, see EX Series Switch Software Features Overview.

NOTE: Nonstop bridging provides a transparent switchover mechanism only for Layer 2 protocol sessions. Nonstop active routing (NSR) provides a similar mechanism for Layer 3 protocol sessions. See "Understanding Nonstop Active Routing on EX Series Switches" on page 9.

**Related Documentation**

- Configuring Nonstop Bridging on EX Series Switches (CLI Procedure) on page 48

 

CHAPTER 4

# Nonstop Software Upgrade Overview



## Understanding Nonstop Software Upgrade on EX Series Switches

Nonstop software upgrade (NSSU) enables you to upgrade the software running on a Juniper Networks EX Series Virtual Chassis or a Juniper Networks EX Series Ethernet Switch with redundant Routing Engines with a single command and minimal disruption to network traffic.

NSSU is supported on the following platforms:

- EX3300 Virtual Chassis
- EX4200 Virtual Chassis
- EX4500 Virtual Chassis
- EX4550 Virtual Chassis
- All Mixed Virtual Chassis composed of EX4200, EX4500, and EX4550 switches
- EX6200 switches
- EX8200 switches
- EX8200 Virtual Chassis

Performing an NSSU provides these benefits:

- No disruption to the control plane—An NSSU takes advantage of graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) to ensure no disruption to the control plane. During the upgrade process, interface, kernel, and routing protocol information is preserved.
- Minimal disruption to traffic—An NSSU minimizes traffic disruption by:
  - Upgrading line cards one at a time in an EX6200 switch, EX8200 switch, or EX8200 Virtual Chassis, permitting traffic to continue to flow through the line cards that are not being upgraded.
  - Upgrading member switches one at a time in an EX3300, EX4200, EX4500, or mixed Virtual Chassis, permitting traffic to continue to flow through the members that are not being upgraded.

To achieve minimal disruption to traffic, you must configure link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards or Virtual Chassis members. When one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.

NOTE: Because NSSU upgrades the software on each line card or on each Virtual Chassis member one at a time, an upgrade using NSSU can take longer than an upgrade using the **request system software add** command.

For EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis, you can reduce the amount of time an upgrade takes by configuring line-card upgrade groups. The line cards 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 (CLI Procedure)" on page 49.

This topic covers:



## Requirements for Performing an NSSU

The following requirements apply to all switches and Virtual Chassis:

- All 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.

  NOTE: Although nonstop bridging (NSB) does not have to be enabled to perform an NSSU, we recommend enabling NSB before performing an NSSU. 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)" on page 48.

- For minimal traffic disruption, you must define link aggregation groups (LAGs) such that the member links reside on different Virtual Chassis members or on different line cards.

The following are requirements for EX3300, EX4200, EX4500, and mixed Virtual Chassis:

- The 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 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 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 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 must have **no-split-detection** configured so that the Virtual Chassis does not split when an NSSU upgrades a member.

## How an NSSU Works

This section describes what happens when you request an NSSU on these switches and Virtual Chassis:


- EX3300, EX4200, EX4500, and Mixed Virtual Chassis on page 15
- EX6200 and EX8200 Switches on page 16
- EX8200 Virtual Chassis on page 17


### EX3300, EX4200, EX4500, and Mixed Virtual Chassis

When you request an NSSU on an EX3300, EX4200, EX4500, or mixed Virtual Chassis:

1. The Virtual Chassis master verifies that:
   - The backup is online and running the same software version.
   - Graceful Routing Engine switchover and nonstop active routing are enabled.
   - The Virtual Chassis has a preprovisioned configuration.
2. The master installs the new software image on the backup and reboots it.
3. The master resynchronizes the backup.
4. The master installs the new software image on member switches that are in the linecard role and reboots them, one at a time. The master waits for each member to become online and active before starting the software upgrade on the next member.
5. 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.
6. 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.

### EX6200 and EX8200 Switches

When you request an NSSU on a standalone switch with redundant Routing Engines:

1. The switch verifies that:
   - Both Routing Engines are online and running the same software version.
   - Both Routing Engines have sufficient storage space for the new software image.
   - Graceful Routing Engine switchover and nonstop active routing are enabled.
2. The switch installs the new software image on the backup Routing Engine and reboots it.
3. The switch resynchronizes the backup Routing Engine to the master Routing Engine.
4. The line cards in the first upgrade group (or the line card in slot 0, if no upgrade groups are defined) download the new image and then restart. Traffic continues to flow through the line cards in the other upgrade groups during this process.
5. When line cards restarted in Step 4 are online again, the line cards in the next upgrade group download the new image and restart. This process continues until all online line cards have restarted with the new software.

   NOTE: **If you have taken a line card offline with the CLI before you start the NSSU, the line card is not restarted and remains offline.**

6. The switch performs a graceful Routing Engine switchover, so that the upgraded backup Routing Engine becomes the master.
7. The switch installs the new software on the original master Routing Engine.

   To complete the upgrade process, the original master Routing Engine must be rebooted. You can do so manually or have the switch perform an automatic reboot by including the **reboot** option when you request the NSSU. After the original master has been rebooted, you can optionally return control to it by requesting a graceful Routing Engine switchover.
8. (EX6200 switch only) The original master Routing Engine reboots to complete the software upgrade.

   NOTE: **To complete the upgrade process on an EX8200 switch, you must intervene to reboot the original master Routing Engine. You can reboot the original master Routing Engine manually or have the switch perform an automatic reboot by including the reboot option when you request the NSSU.**

9. (Optional) After the original master has been rebooted, you can return control to it by requesting a graceful Routing Engine switchover.

The switch can maintain normal operations with either Routing Engine acting as the master Routing Engine after the software upgrade, so you only have to perform this switchover if you want to return Routing Engine control to the original master Routing Engine.

### EX8200 Virtual Chassis

When you request an NSSU on an EX8200 Virtual Chassis:

1. The master external Routing Engine verifies that:
   - It has a backup external Routing Engine that is online.
   - All Virtual Chassis members have redundant Routing Engines and the Routing Engines are online.
   - All Routing Engines are running the same software version.
   - All Routing Engines have sufficient storage space for the new software image.
   - Graceful Routing Engine switchover and nonstop active routing (NSR) are enabled.
2. The master external Routing Engine installs the new software image on the backup external Routing Engine and reboots it.
3. The backup external Routing Engine resynchronizes with the master external Routing Engine.
4. The master external Routing Engine installs the new software on the backup Routing Engines in the member switches and reboots the backup Routing Engines.
5. When the reboot of the backup Routing Engines complete, the line cards in the first upgrade group download the new image and then restart. (If no upgrade groups are defined, the line card in slot 0 of member 0 downloads the new image and restarts.) Traffic continues to flow through the line cards in the other upgrade groups during this process.
6. When line cards restarted in Step 5 are online again, the line cards in the next upgrade group (or the next sequential line card) download the new image and restart. This process continues until all online line cards have restarted with the new software.

NOTE: **If you have taken a line card offline with the CLI before you start the NSSU, the line card is not restarted and remains offline.**

7. The new software image is installed on the master Routing Engines, both external and internal.
8. The member switches perform a graceful Routing Engine switchover, so that the upgraded backup Routing Engines become masters.
9. The master external Routing Engine performs a graceful Routing Engine switchover so that the backup external Routing Engine is now the master.

To complete the upgrade process, the original master Routing Engines, both external and internal, must be rebooted. You can do so manually by establishing a console

connection to each Routing Engine or have the reboot performed automatically by including the **reboot** option when you request the NSSU. After the original master external Routing Engine has been rebooted, 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 rollback 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 must be running a Junos OS release that supports NSSU before you can perform an NSSU. If a Virtual Chassis is running a software version that does not support NSSU, use the **request system software add** command.

Table 3 on page 18 lists the EX Series switches and Virtual Chassis that support NSSU and the Junos OS release at which they began supporting it.

### Table 3: Platform and Release Support for NSSU

| Platform | Junos OS Release |
|---|---|
| EX3300 Virtual Chassis | 12.2 or later |
| EX4200 Virtual Chassis | 12.1 or later |
| EX4500 Virtual Chassis | 12.1 or later |
| EX4550 Virtual Chassis | 12.2 or later |
| Mixed EX4200 and EX4500 Virtual Chassis | 12.1 or later |
| Mixed EX4200 and EX4550 Virtual Chassis | 12.2 or later |
| Mixed EX4200, EX4500, and EX4550 Virtual Chassis | 12.2 or later |
| Mixed EX4500 and EX4550 Virtual Chassis | 12.2 or later |
| EX6200 switch | 12.2 or later |
| EX8200 switch | 10.4 or later |
| EX8200 Virtual Chassis | 11.1 or later |

## 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 14. NSSU requires no additional configuration.

For EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis, you can optionally configure line-card upgrade groups using the CLI. See "Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches" on page 44.

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.

**Related Documentation**

- Upgrading Software on an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, or Mixed Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) on page 108
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure) on page 95
- Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) on page 103
- Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure) on page 47
- Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)
- Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 44

 

CHAPTER 5

# Power Management Overview



## Understanding Power Management on EX Series Switches

The power management feature for Juniper Networks EX6200 Ethernet Switches and Juniper Networks EX8200 Ethernet Switches helps ensure that normal operation of the system is not disrupted because of insufficient power to the switch. For example:

- Power management ensures that operating line cards continue to receive power if a user installs a new line card in an operating switch when power is insufficient for both the new and existing line cards.
- Power management reserves a certain amount of power to power supply redundancy, so that if a power supply fails, the switch can continue to operate normally. If power management must use some of this reserved power to provide power to switch components, it raises an alarm to indicate that power supply redundancy no longer exists and that normal operations might be disrupted if a power supply fails.
- If power supply failure requires power management to power down some components, it does so gracefully by powering down line cards and PoE ports in the order specified by the user.

Power management manages power to switch components by employing a power budget policy. In its power budget policy, power management:

- Budgets power for each installed switch component that requires power. With the exception of PoE power for line cards that support PoE, the amount that power management budgets for each component is the maximum power that component might consume under worst case operating conditions. For example, for the fan tray, power management budgets the amount of power required to run the fans at their maximum speed setting, even if the current fan speed is much lower.
- Reserves a set amount of power for power supply redundancy. In its default configuration, power management manages the switch for N+1 power redundancy, which ensures uninterrupted system operation if one power supply fails. For example, if a switch has four online 3000 W power supplies, power management reserves 3000 W in its power budget policy for redundancy. It allocates the remaining 9000 W to normal operating power.

- Specifies the rules under which components receive power. These rules are designed to ensure the least disruption to switch operation under conditions of insufficient power. For example, power management provides power to core system components, such as the Routing Engines, before it provides power to line cards.

You can configure certain aspects of power management's budget policy, specifically:

- The power priority of individual line cards. By assigning different power priorities to the line cards, you can determine which line cards are more likely to receive power in the event of insufficient power.
- The power redundancy configuration. The default power redundancy configuration is N+1; you can optionally configure N+N. For example, if you have deployed two independent AC power feeds to the switch, configure N+N redundancy. When you configure power management for N+N redundancy, it reserves the appropriate amount of power in its power budget and reports insufficient power conditions accordingly.

These configurable items are discussed further in:


- Power Priority of Line Cards on page 22
- Power Supply Redundancy on page 25


## Power Priority of Line Cards

The power priority of line cards determines:

- The order in which line cards are allocated power
- The order in which line cards that support PoE are allocated power for PoE
- How power is reallocated in cases of changes in power availability or demand in an operating switch

NOTE: **On EX6200 switches, the four 10-Gigabit Ethernet SFP+ uplink ports on a Switch Fabric and Routing Engine (SRE) module are treated like a line card in the power budget.**

This section covers:


- How a Line Card's Power Priority Is Determined on page 22
- Line Card Priority and Line Card Power on page 23
- Line Card Priority and PoE Power on page 23
- Line Card Priority and Changes in the Power Budget on page 24


### How a Line Card's Power Priority Is Determined

Using the CLI, you can assign a explicit power priority to a line-card slot. If more than one slot has the same assigned priority, the power priority is determined by slot number, with the lowest-numbered slots receiving power first.

By default, all slots in an EX8200 switch are assigned the lowest priority. Thus if you do not explicitly assign priorities to slots, power priority is determined by slot number, with slot 0 having the highest priority.

In an EX6200 switch, all slots are assigned the lowest priority, except for the slots containing an SRE module. Slots containing an SRE module are automatically assigned the highest priority. This means that the line cards that represent the 10-Gigabit Ethernet SFP+ ports on SRE modules have the highest priority among the line cards.

### Line Card Priority and Line Card Power

When an EX6200 or EX8200 switch is powered on, power management allocates power to components according to its power budget policy. After power management has allocated power to the base chassis components, it allocates the remaining available power to the line cards. It powers on the line cards in priority order until all line cards are powered on or the available power (including reserved power, if necessary) is exhausted. Thus if available power is exhausted before all line cards receive power, higher-priority cards are powered on while lower-priority cards remain powered off.

A lower-priority card might receive power while a higher-priority card does not if the remaining available power is sufficient to power on the lower-priority card but not the higher-priority card. For example, if a line card requiring 450 W is in a higher-priority slot than line card requiring 330 W, the line card requiring 330 W receives the power if there is less than 450 W but more than 330 W remaining in the power budget.

Line cards that have been administratively taken offline are not allocated power.

NOTE: Because power management does not allocate power to a line card that has been administratively taken offline, a line card that has been taken offline in an EX6200 or EX8200 switch is not automatically brought online when you commit a configuration. You must explicitly use the request chassis fpc slot *slot-number* online command to bring a line card online that was taken offline previously. This behavior differs from other platforms running Juniper Networks Junos operating system (Junos OS), which automatically bring an offline FPC online when you commit a configuration.

If power management cannot power on a line card because of insufficient power, it raises a major (red) alarm.

### Line Card Priority and PoE Power

After all line cards have been powered on, power management allocates any remaining available power, including reserved power, to the PoE power budgets of line cards that have PoE ports. Power management allocates PoE power to line cards in the order of power priority. If enough power is available, a line card receives its full PoE power budget before power management allocates PoE power to the next highest-priority line card. If not enough power is available, a line card receives partial PoE power and lower-priority line cards receive no PoE power.

If power management is unable to allocate enough power to meet the PoE power budget for a line card, it logs a message to the system log.

The default PoE power budget for a line card is the amount of power needed to supply the maximum supported power to all PoE ports. In cases where powered devices do not require the maximum power or in which some PoE ports are not used for powered devices, you can configure a smaller PoE power budget for a line card. By configuring a smaller PoE power budget, you make more power available for the PoE power budgets of lower-priority line cards.

You can also configure the power priority of the PoE ports on a line card. If power management is unable to allocate enough power to a line card to meet its PoE power budget, the line card PoE controller will turn off power to PoE ports in reverse priority order as required to meet the reduced power allocation.

See Configuring PoE (CLI Procedure) for more information on how to configure the PoE power budget for a line card and how to configure PoE port priorities.

### Line Card Priority and Changes in the Power Budget

In an operating switch, power management dynamically reallocates power in response to changes in power availability or demand or changes in line card priority. Power management uses line card priority to determine how to reallocate power in response to the following events:

- A power supply fails, is removed, or is taken offline:
  - If power is insufficient to meet the PoE power allocations of all PoE line cards, power management deallocates PoE power from the line cards in reverse priority order until power is sufficient to meet the remaining PoE power allocations.
  - If power is insufficient to meet the base (non-PoE) power requirements of all the line cards, all PoE power is deallocated. If, after the deallocation of PoE power, power is still not sufficient, power management turns off line cards in reverse priority order until power is sufficient for the remaining line cards.
- A new line card is inserted or a line card is brought online:
  - If the line card supports PoE and there is insufficient power to meet its PoE power budget, PoE power is reallocated from lower-priority line cards. If not enough PoE power can be reallocated from lower-priority line cards, the new line card receives a partial PoE power allocation.
  - If there is insufficient power to power on the new line card, PoE power is removed from PoE line cards in reverse priority order until the new line card can be powered on.
  - If the removal of all PoE power is insufficient to free up enough power to power on the line card, the line card remains powered off and the PoE line cards continue to receive their PoE power allocations. To minimize disruption on an operating switch, lower-priority line cards are not turned off to provide power to the new line card. However, if you restart the switch, power management reruns the current power budget policy and powers line cards on or off based on their priority. As a result, line

cards receive power strictly by priority order and previously operating line cards might no longer receive power.

- A new power supply is brought online:
  - Any line cards that were powered off because of insufficient power are powered on in priority order.
  - After all line cards are powered on, remaining power is allocated to the PoE power budgets of line cards in priority order.
- A line card is removed or taken offline, freeing up power:
  - Any line cards that were powered down because of insufficient power are powered on in priority order.
  - After all line cards are powered on, any remaining power is allocated to the PoE power budgets of line cards in priority order.
- A user changes the assigned power priority of one or more line cards when power is insufficient to meet the power budget:
  - PoE power to the line cards is reallocated based on the new power priorities.
  - Base power allocation to the line cards is not changed—in other words, power management does not power down line cards that had been receiving power because they are now a lower priority. However, if you restart the switch, power management reruns the current power budget policy and powers line cards on or off based on their priority. As a result, line cards receive power strictly by priority order and previously operating line cards might no longer receive power.

If, because of insufficient power, power management reduces or eliminates the PoE power budget for a line card, it logs a message to the system log. If power management must power down a line card because of insufficient power, it raises a major (red) alarm.

## Power Supply Redundancy

By default, power management in EX6200 and EX8200 switches is configured to manage the power supplies for N+1 redundancy, in which one power supply is held in reserve for backup if one of the other power supplies is removed or fails.

You can configure power management to manage the power supplies for N+N redundancy. In N+N redundancy, power management holds N power supplies in reserve for backup. For example, if your switch has six power supplies and you configure N+N redundancy, power management makes three power supplies available for normal operating power and reserves three power supplies for redundancy (3+3). If you have an odd number of power supplies, power management allocates one more power supply to normal operating power than to redundant power. For example, if you have five power supplies, the N+N configuration is 3+2.

Given the same number of power supplies, an N+N configuration usually provides less normal operating power than an N+1 configuration because the N+N configuration holds more power in reserve for backup. Table 4 on page 26 shows the effect on normal operating power in N+1 and N+N configurations.

 

### Table 4: Available Operating Power in N+1 and N+N Redundancy Configurations

| Number of Power Supplies at *n* W Each | Normal Operating Power in N+1 Configuration | Normal Operating Power in N+N Configuration |
|---|---|---|
| 2 | 1 x (*n* W) | 1 x (*n* W) |
| 3 | 2 x (*n* W) | 2 x (*n* W) |
| 4 | 3 x (*n* W) | 2 x (*n* W) |
| 5 (EX8200 switches only) | 4 x (*n* W) | 3 x (*n* W) |
| 6 (EX8200 switches only) | 5 x (*n* W) | 3 x (*n* W) |

To compensate for the reduced normal operating power, power management on EX8200 switches allocates less power to the chassis in an N+N configuration than in an N+1 configuration. This reduction in allocated chassis power allows a switch in an N+N configuration to power more line cards than it could without the reduction. For the EX8208 switch, the power allocated for the chassis is reduced to 1200 W from 1600 W; for the EX8216 switch, it is reduced to 1800 W from 2400 W.

NOTE: To achieve the reduction in allocated chassis power in an EX8200 switch, power management reduces the maximum fan speed to 60 percent in an N+N configuration from 80 percent in an N+1 configuration. Because the maximum fan speed is reduced, it is possible that a line card that overheats would be shut down sooner in an N+N configuration than in an N+1 configuration.

On EX6200 switches, the same amount of power is allocated for the chassis in N+N configurations as in N+1 configurations.

Power management automatically recalculates the reserved power and normal operating power as power supplies go online or offline. For example, if you have an N+N configuration with three online 2000 W power supplies, power management allocates 2000 W to reserved power. If you bring a fourth 2000 W power supply online, power management then allocates 4000 W to reserved power. If a power supply goes offline again, power management once again allocates 2000 W to reserved power.

When power is insufficient to meet the budgeted power requirements, power management raises alarms as follows:

- A minor (yellow) alarm is raised when insufficient power exists to maintain the configured N+1 or N+N power reserves, but all line cards are still receiving their base and PoE power allocations. If this condition persists for 5 minutes, the alarm becomes a major (red) alarm. Even though operation of the switch is unaffected in this condition, you should remedy it as quickly as possible because a power supply failure might cause a disruption in switch operation.

- A major (red) alarm is raised when insufficient power exists to provide all the line cards with their base and PoE power allocations. One or more PoE ports might be down or one or more line cards might be down.

Power management clears all alarms when sufficient power is available to meet normal operating and reserved power requirements.

**Related Documentation**

- Understanding Alarm Types and Severity Levels on EX Series Switches
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 53
- Configuring Power Supply Redundancy (CLI Procedure) on page 51
- Verifying Power Configuration and Use on page 113

 

CHAPTER 6

# Redundant Power System Overview



## EX Series Redundant Power System Hardware Overview

The Redundant Power System (RPS) can be used to provide backup power for Juniper Networks EX2200 Ethernet Switches, (except Juniper Networks EX2200-C Ethernet Switches) and Juniper Networks EX3300 Ethernet Switches that are standalone switches or are members of a Virtual Chassis.

Most EX Series switches have a built-in capability for redundant power supplies—therefore, if one power supply fails on those switches, the other power supply takes over. However, EX2200 switches and EX3300 switches have only one internal fixed power supply. If an EX2200 switch or EX3300 switch is deployed in a critical situation, we recommend that you connect a an RPS to that switch to supply backup power during a loss of power.

RPS is not a primary power supply—it only provides backup power to switches when the single dedicated power supply fails. An RPS operates in parallel with the single dedicated power supplies of the switches connected to it and provides all connected switches enough power to support either Power over Ethernet (PoE) or non-PoE devices when the power supplies on the switches fail..

An RPS can hold up to three power supplies connected to as many as six switches—how that power is allocated is up to you. You determine whether or not to connect switches that provide PoE and you determine which switches have priority. Priority becomes an issue when you connect more than three switches that provide PoE to a fully loaded RPS because a switch providing PoE requires more power than a switch that does not provide PoE. Because a power supply can support only one switch providing PoE, the RPS can become oversubscribed when too many switches that must have enough power for PoE have a power failure.

This topic describes:

## Switch Models and Configurations Supported by the RPS

The RPS supports all EX3300 switches and EX2200 switches except EX2200-C switches. You can simultaneously connect any supported switches to the same RPS, whether the switches are standalone switches or are configured in a Virtual Chassis.

All power provided by RPS is either PoE or non-PoE. By default, RPS supports switches that provide PoE. If even one switch provides PoE, then the RPS must be configured to provide enough power for PoE. When enough power for PoE is supplied, one switch can be powered by each power supply. If the switches are not providing PoE power, two switches can be powered by one RPS power supply—you can reconfigure an RPS to provide non-PoE power using a feature called multi-backup.

Table 5 on page 30 lists some possible scenarios and RPS solutions. These examples assume that each RPS is fully loaded with three power supplies.

### Table 5: Sample Requirements and RPS Solutions

| Switches Requiring Backup | You need this RPS configuration: |
|---|---|
| Six switches that do not provide PoE to attached devices | One RPS can simultaneously provide power to all six switches if you change the power default to multi-backup—this indicates that no attached switch provides PoE to any devices. |
| One switch that provides PoE to other devices or two switches that do not provide PoE to any devices | One RPS will always back up all three switches, whether or not they provide PoE to connected devices. Leave the power at the default setting (no multi-backup) and let RPS determine that two switches need only minimum power and one switch provides PoE and therefore needs extra power. RPS automatically supplies the correct level of power. |
| One EX Series Virtual Chassis member that supplies PoE, one switch that supplies PoE, and one switch that does not supply PoE to any connected devices | One RPS will always back up all three switches. Leave the power default setting (no multi-backup) and let RPS determine that one switch needs only minimum power, one switch needs extra power because it supplies PoE, and the Virtual Chassis member also provides PoE to connected devices. |
| One switch that supplies PoE and five switches that do not supply PoE | You have two options.<br><br>Option 1—Use one RPS: Up to three switches that do or do not supply PoE can be backed up simultaneously. You can prioritize the six switches to determine which three are most important if all six fail at once. You must leave the power default setting (no multi-backup) because you have one switch that supplies PoE to attached devices and therefore requires more power.<br><br>Option 2—Use Two RPSs: In this case, you can connect three switches to each RPS and all switches will be backed up if they all fail at once. Alternatively, you can change the power default to multi-backup on one RPS and connect all five switches that do not supply PoE to that RPS, leaving the other RPS to back up the switch that supplies PoE. |
| EX Series Virtual Chassis | Use as many RPSs as needed to back up all members of the Virtual Chassis. |

## When a Switch's Power Supply Fails

Because the power supplies for both EX3300 switches and EX2200 switches are internal, if the switch's power supply fails, you must replace the switch. You should remove or replace a switch with a failed power supply as soon as possible.

Do not try to use an RPS as a primary power supply because an RPS cannot boot or reboot a switch. Each switch connected to the RPS must have its own dedicated power supply and must have booted up using the internal power supply.

If a switch is deployed in a large network center where RPS has a separate source of electricity than the switches it supports, the RPS supplies power when only the switch's electricity fails. In this case, you would not have to replace the switch because the power supply is still functional. An EX3300 switch will resume using its own internal power supply when electricity to the switch is restored. However, an EX2200 switch operates differently—once the RPS is supplying power, the switch does not automatically resume using its own internal power supply when power is restored to it.

## Components of the RPS

Table 6 on page 31 lists and describes the components of an RPS:

### Table 6: Redundant Power System Components

| Component | Value |
|---|---|
| Power supplies that can be installed | Up to three EX-PWR3-930-AC power supplies. One is included and additional power supplies must be ordered separately. |
| Switch connector ports on RPS | 6 (2 per power supply) |
| Power cords<br>(for connecting power supplies to the AC power source outlet) | Up to three power cords, one per power supply. |
| RPS cables<br>(for connecting a switch to a power supply installed in the RPS) | 6 (1 for each RPS-to-switch connection). One cable is supplied with the RPS. Additional cables must be ordered separately. |

**Related Documentation**

- Physical Description of a Redundant Power System
- Chassis Physical Specifications for the EX Series Redundant Power System
- Determining and Setting Priority for Switches Connected to an EX Series RPS on page 54
- Installing a Power Supply in the EX Series Redundant Power System

## Understanding How Power Priority Is Determined and Set for Switches Connected to the EX Series Redundant Power System

The Redundant Power System (RPS) is designed to provide backup power to switches that lack built-in redundant power supplies. The RPS provides backup power to switches that either supply power over Ethernet (PoE), which require more power, or switches that do not supply PoE, which require less power. A power supply can either power one PoE device or two non-PoE devices. That means if an RPS is fully loaded with three power supplies, supports PoE switches, and more than three PoE switches have a power failure, some switches will not be powered. You can, however, determine which switches will be powered when an RPS is oversubscribed. When too many connected switches fail, the switches are given power based on their priority. Priority is also reconfigured when any power change takes place. For example, if three switches are already being backed up and another switch has a power failure, the RPS detects this, reconfigures the current top priorities, and allots power accordingly.


- Default RPS Priority on page 32
- Changing the Priority of Switches on an EX Series RPS on page 32


### Default RPS Priority

While six non-PoE switches can all simultaneously be backed up with three power supplies, only three PoE switches can be backed up (because PoE uses more power). This means that an RPS with four or more PoE switches connected will have to select three of them for backup. You can determine priority by the connector positions you use to connect the switches. By default, an RPS assigns priority to switches based on their switch connector port location, with the leftmost port having the lowest priority and the rightmost port having the highest priority. If the PoE switches shown in Figure 1 on page 32 all fail, the manufacturing, support, and finance switches will be backed up because they are connected to the rightmost connectors.

**Figure 1: Default PoE Switch Priority Is Determined by Connector Port Location**

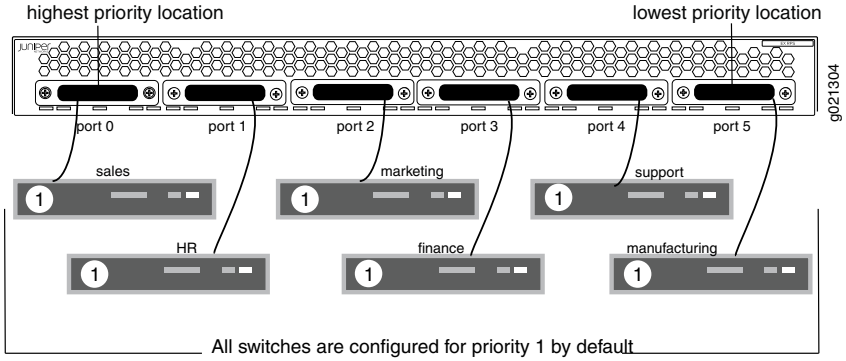


### Changing the Priority of Switches on an EX Series RPS

There is a way to alter the priority of PoE switches on an RPS without disconnecting the cables. You can optionally reconfigure any of the attached switches from their CLIs to

establish a switch's RPS priority—this CLI configuration overcomes the priority determined by the switch connector port location. Priority ranges from zero (off) to 1 (lowest) through 6 (highest). By default, all switches are configured to 1, the lowest priority. Let's say that the sales switch is reconfigured from the switch's CLI for priority 5 (second highest).

Now in Figure 2 on page 33, with the sales switch configured for RPS 5 from the CLI, the highest priority changes to sales (because 5 is higher than 1), then manufacturing, and then support.

### Figure 2: Switch Priority After CLI Configuration

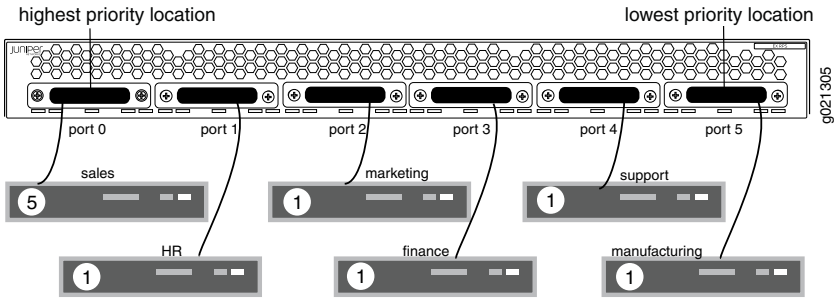


When assigning power priority to switches by using the CLI on the switch, keep these points in mind:

- By default, all switches are assigned priority 1 (lowest) and derive precedence from the location of their connector port on the RPS, with the rightmost port having highest priority.
- Priority 0 assigned from a switch CLI means that the RPS does not provide any backup power to the switch. Essentially, this turns off RPS support.
- Priority 6 assigned from a switch CLI is the highest priority and priority 1 is the lowest priority.
- The CLI command that assigns priority to EX2200 switches is slightly different from the CLI command that assigns priority to EX3300 switches because EX3300 switches can be configured as a Virtual Chassis.
- If two or more switches are assigned the same priority value from the switches' CLIs , then the power priority for those switches is determined by the RPS switch connector port location, with the ports to the right receiving priority.
- If a single power supply is installed, the RPS can provide backup power to one switch out of all the switches connected to the RPS. If you do not need any PoE power backup on any switch, you can increase the number of supported switches to two per power supply. Switches connected to an RPS must be either all PoE or all non-PoE.
- The RPS discontinues supplying backup power to a lower-priority switch if it detects a backup power need for a higher-priority switch at the same time.

**Related Documentation**

- Determining and Setting Priority for Switches Connected to an EX Series RPS on page 54
- Setting Priority for Switches Connected to the Redundant Power System (CLI Procedure)

 

CHAPTER 7

# VRRP Overview



## Understanding VRRP on EX Series Switches

Juniper Networks EX Series Ethernet Switches support the Virtual Router Redundancy Protocol (VRRP) and VRRP for IPv6. This topic covers:



### Overview of VRRP on EX Series Switches

You can configure the Virtual Router Redundancy Protocol (VRRP) or VRRP for IPv6 on Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and logical interfaces on EX Series switches. When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup EX Series switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts. Virtual Router Redundancy Protocol is not supported on management interfaces.

VRRP for IPv6 provides a much faster switchover to an alternate default routing platform than IPv6 Neighbor Discovery (ND) procedures. VRRP for IPv6 does not support the **authentication-type** or **authentication-key** statements.

NOTE: Do not confuse the VRRP master and backup routing platforms with the master and backup member switches of a Virtual Chassis configuration. The master and backup members of a Virtual Chassis configuration compose a single host. In a VRRP topology, one host operates as the master routing platform and another operates as the backup routing platform, as shown in Figure 4 on page 37.

Switches running VRRP dynamically elect master and backup routing platforms. You can also force assignment of master and backup routing platforms using priorities from **1** through **255**, with **255** being the highest priority. In VRRP operation, the default master routing platform sends advertisements to backup routing platforms at regular intervals. The default interval is 1 second. If the backup routing platforms do not receive an advertisement for a set period, the backup routing platform with the highest priority takes over as master and begins forwarding packets.

NOTE: Priority 255 cannot be set for routed VLAN interfaces (RVIs).

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

## Examples of VRRP Topologies

Figure 3 on page 36 illustrates a basic VRRP topology with EX Series switches. In this example, Switches A, B, and C are running VRRP and together they make up a virtual routing platform. The IP address of this virtual routing platform is **10.10.0.1** (the same address as the physical interface of Switch A).

### Figure 3: Basic VRRP on EX Series Switches

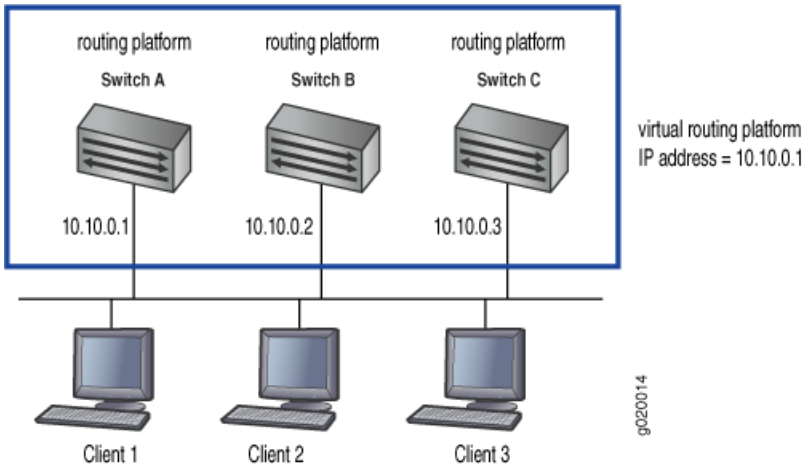


Figure 4 on page 37 illustrates a basic VRRP topology using Virtual Chassis configurations. Switch A, Switch B, and Switch C are each composed of multiple interconnected Juniper Networks EX4200 Ethernet Switches. Each Virtual Chassis configuration operates as a single switch, which is running VRRP, and together they make up a virtual routing platform.

The IP address of this virtual routing platform is **10.10.0.1** (the same address as the physical interface of Switch A).

Figure 4: VRRP on Virtual Chassis Switches

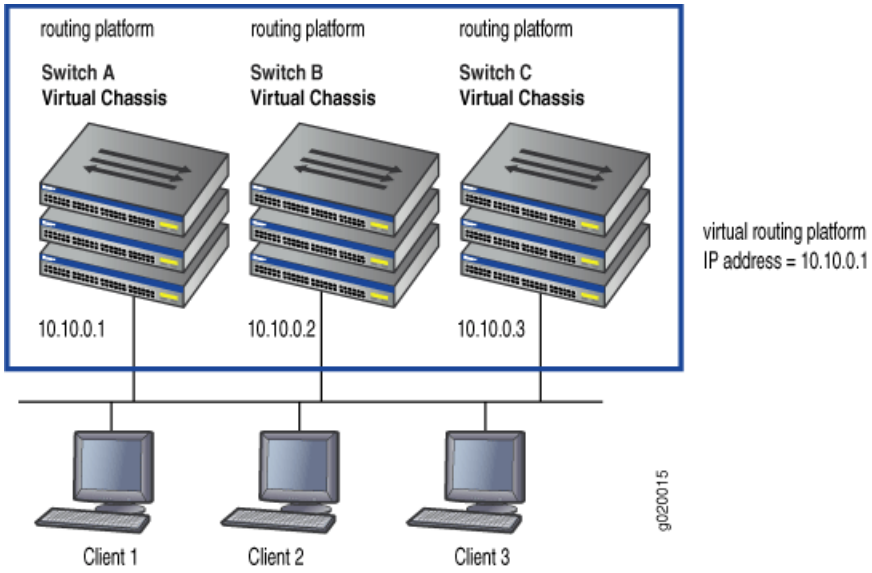


Because the virtual routing platform uses the IP address of the physical interface of Switch A, Switch A is the master VRRP routing platform, while Switch B and Switch C function as backup VRRP routing platforms. Clients 1 through 3 are configured with the default gateway IP address of **10.10.0.1** as the master router, Switch A, forwards packets sent to its IP address. If the master routing platform fails, the switch configured with the higher priority becomes the master virtual routing platform and provides uninterrupted service for the LAN hosts. When Switch A recovers, it becomes the master virtual routing platform again.

**Related Documentation**

- For more information on VRRP or VRRP for IPv6, see the *Junos OS High Availability Configuration Guide*.
- High Availability Features for EX Series Switches Overview on page 3
- Configuring VRRP for IPv6 (CLI Procedure) on page 53

 

PART 2

# Configuration

 

CHAPTER 8

# Configuration Examples



## Example: Configuring Nonstop Active Routing on EX Series Switches

Nonstop active routing (NSR) provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

This example describes how to configure nonstop active routing on switches with multiple Routing Engines:



### Requirements

This example uses the following hardware and software components:

- An EX Series switch with multiple Routing Engines
- Junos OS Release 10.4 or later for EX Series switches

### Overview and Topology

Configure nonstop active routing on any EX Series switch with multiple Routing Engines or EX Series switch in a Virtual Chassis configuration. Nonstop active routing is advantageous in networks where neighbor routing devices do not support graceful restart protocol extensions.

The topology used in this example consists of an EX8200 switch with redundant Routing Engines connected to neighbor routing devices that are not configured to support graceful restart of protocols.

## Configuration

**CLI Quick Configuration**

To quickly configure nonstop active routing, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis redundancy graceful-switchover
set routing-options nonstop-routing
set system commit synchronize
```

**Step-by-Step Procedure**

To configure nonstop active routing on a switch:

1. Enable graceful Routing Engine switchover (GRES):

   ```
   [edit chassis redundancy]
   user@switch# set graceful-switchover
   ```

2. Enable nonstop active routing (by default, nonstop active routing is disabled):

   ```
   [edit routing-options]
   user@switch# set nonstop-routing
   ```

3. Synchronize configuration changes between the Routing Engines:

   ```
   [edit system]
   user@switch# set commit synchronize
   ```

   If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.

NOTE: If the backup Routing Engine is down when you issue the commit, a warning is displayed and the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes up, its configuration is automatically synchronized with that of the master. If you subsequently insert or bring up a backup Routing Engine, it automatically synchronizes its configuration with the master Routing Engine configuration.

### Results

Check the results of the configuration:

```
[edit]
user@switch# show
chassis {
    redundancy {
        graceful-switchover;
    }
routing-options {
    nonstop-routing;
}
system {
```

```
    commit synchronize;
}
```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Nonstop Active Routing Is Working Correctly on the Switch on page 43

### Verifying That Nonstop Active Routing Is Working Correctly on the Switch

**Purpose** Verify that nonstop active routing is enabled.

**Action** Issue the **show task replication** command:

```
user@switch# show task replication
Stateful Replication: Enabled
RE mode: Master

 Protocol                Synchronization Status
 OSPF                    Complete
 RIP                     Complete
 PIM                     Complete
 RSVP                    Complete
```

**Meaning** This output shows that nonstop active routing (Stateful Replication) is enabled on master routing engine. If nonstop routing is not enabled, instead of the output shown above:

- On the backup routing engine the following error message is displayed: "**error: the routing subsystem is not running**."
- On the master routing engine, the following output is displayed if nonstop routing is not enabled:

  ```
  Stateful Replication: Disabled
  RE mode: Master
  ```

## Troubleshooting

To troubleshoot nonstop active routing, perform these tasks:

- Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled on page 43

### Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled

**Problem** A protocol loses connectivity with neighbors after a graceful Routing Engine switchover (GRES) occurs with nonstop active routing (NSR) enabled.

**Solution** Use trace options to help isolate the problem and gather troubleshooting information. Using the information gathered from trace options, you can confirm or eliminate the synchronization of the Routing Engines as the cause of the loss of connectivity for the protocol. See "Tracing Nonstop Active Routing Synchronization Events" on page 159.

 

**Related Documentation**

- Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure) on page 47
- Tracing Nonstop Active Routing Synchronization Events on page 159
- Understanding Nonstop Active Routing on EX Series Switches on page 9

## Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Nonstop software upgrade (NSSU) enables you to upgrade the software running on an EX Series switch with redundant Routing Engines or on most EX Series Virtual Chassis by using a single command and with minimal disruption to network traffic. By default, NSSU upgrades the software running on line cards one line card at a time.

To reduce the time an NSSU takes, you can configure line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines or on an EX8200 Virtual Chassis.

NOTE: **NSSU line-card upgrade groups are not supported for NSSUs on EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, or any mixed Virtual Chassis composed of EX4200, EX4500, or EX4550 switches.**

This example shows how to configure NSSU to use line-card upgrade groups:

- Requirements on page 44
- Overview and Topology on page 44
- Configuration on page 45

### Requirements

This example uses the following hardware and software components:

- An EX8200 switch with redundant Routing Engines
- Junos OS Release 10.4 or later for EX Series switches

Before you begin to configure line-card upgrade groups, ensure that you have configured the link aggregation groups (LAGs) as described in Configuring Aggregated Ethernet Links (CLI Procedure). See "Overview and Topology" on page 44 for details about the LAG configurations for this example.

### Overview and Topology

In its default configuration, NSSU upgrades each line card in a switch or Virtual Chassis one at a time. Traffic continues to flow through the other line cards while a line card is being restarted as part of the upgrade. This behavior allows you minimize disruption to traffic by configuring link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards. When one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.

Because the default configuration upgrades each line card one at a time, the upgrade can take some time to complete. You can reduce the time it takes to perform an NSSU by configuring line-card upgrade groups. Instead of being upgraded sequentially, the line cards in an upgrade group are upgraded simultaneously. 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 that are in different upgrade groups.

This example uses an EX8200 switch that has five line cards installed in slots 0 through 4. Two LAGs have been configured:

- **ae0**—Has two member links, one on the line card in slot 0 and one on the line card in slot 1.
- **ae1**—Has two member links, one on the line card in slot 2 and one on the line card in slot 3.

The interfaces on the line card in slot 4 are not part of either LAG.

To minimize the time an upgrade takes and to ensure that the member links of each LAG are in different upgrade groups, this example configures the following two line-card upgrade groups:

- **group1**—Contains the line cards in slots 0, 2, and 4.
- **group2**—Contains the line cards in slots 1 and 3.

The line card in slot 4 could be put in either group. It could also be left out of an upgrade group entirely, and it would be upgraded separately after the line cards in the upgrade groups have been upgraded. However, it is more efficient to include it in an upgrade group.

Figure 5 on page 45 illustrates the topology.

Figure 5: Example Line-Card Upgrade Group Topology

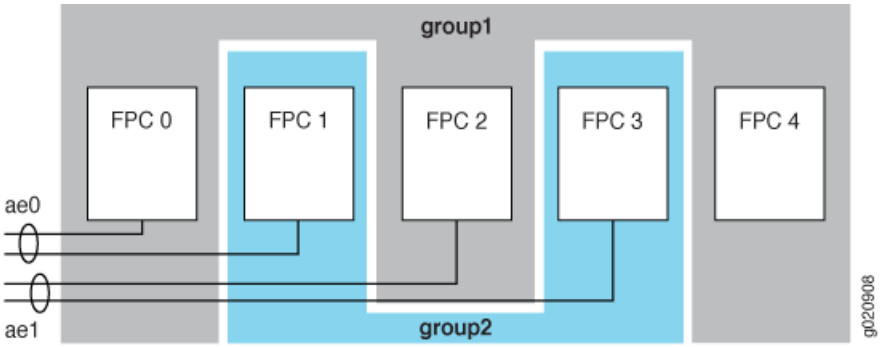


## Configuration

To create line-card upgrade groups, perform these tasks:

**CLI Quick Configuration**

To quickly create the line-card upgrade groups, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis nssu upgrade-group group1 fpcs [0 2 4]
```

**set chassis nssu upgrade-group group2 fpcs [1 3]**

**Step-by-Step Procedure**

To create the line-card upgrade groups for an NSSU:

1. Create the first line-card upgrade group:

```
[edit chassis]
user@switch# set nssu upgrade-group group1 fpcs [0 2 4]
```

2. Create the second line-card upgrade group:

```
[edit chassis]
user@switch# set nssu upgrade-group group2 fpcs (NSSU Upgrade Groups) [1 3]
```

## Results

**Results**

Display the results of the configuration:

```
[edit chassis]
user@switch# show
nssu {
    upgrade-group group1 {
        fpcs [ 0 2 4 ];
    }
    upgrade-group group2 {
        fpcs [ 1 3 ];
    }
}
```

**Related Documentation**

- Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure) on page 49
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure) on page 95
- Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) on page 103

CHAPTER 9

# Configuration Tasks



## Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure)

Nonstop active routing (NSR) provides a mechanism for transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

You can configure nonstop active routing on an EX Series switch with redundant Routing Engines to enable the transparent switchover of the Routing Engines in the event that the Routing Engines switch over. You can also disable nonstop active routing after you have enabled it.

To configure nonstop active routing:

1. Enable graceful Routing Engine switchover (GRES):

   ```
   [edit chassis redundancy]
   user@switch# set graceful-switchover
   ```

2. Enable nonstop active routing (by default, nonstop active routing is disabled):

   ```
   [edit routing-options]
   user@switch# set nonstop-routing
   ```

3. Synchronize configuration changes between the Routing Engines:

   ```
   [edit system]
   user@switch# set commit synchronize
   ```

   If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.

NOTE: There is no requirement to start the two Routing Engines simultaneously. If the backup Routing Engine is not up when you **commit synchronize**, the candidate configuration is committed in the master Routing Engine and when the backup Routing Engine is inserted or comes online, its configuration is automatically synchronized with that of the master.

BEST PRACTICE: After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics** (*interface-name* | **all**) command to reset the cumulative values for local statistics on the new master Routing Engine.

To disable nonstop active routing:

```
[edit routing-options]
user@switch# delete nonstop-routing
```

**Related Documentation**

- Example: Configuring Nonstop Active Routing on EX Series Switches on page 41
- Tracing Nonstop Active Routing Synchronization Events on page 159
- Understanding Nonstop Active Routing on EX Series Switches on page 9

## Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on an EX Series switch with redundant Routing Engines.

Nonstop bridging operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because they are already synchronized on the backup Routing Engine. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the Layer 2 protocol sessions on the switch.

To configure nonstop bridging:

1. Enable graceful Routing Engine switchover (GRES):

   ```
   [edit chassis redundancy]
   user@switch# set graceful-switchover
   ```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

   ```
   [edit system]
   user@switch# set commit synchronize
   ```

   If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit ethernet-switching-options]
user@switch# set nonstop-bridging
```

NOTE: There is no requirement to start both Routing Engines simultaneously. If the backup Routing Engine is not up when you commit the configuration, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes online, the configuration is automatically synchronized.

**Related Documentation**

- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on EX Series Switches
- Understanding Nonstop Bridging on EX Series Switches on page 11

## Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure)

Nonstop software upgrade (NSSU) enables you to upgrade the software running on an EX Series switch with redundant Routing Engines or on most EX Series Virtual Chassis by using a single command and with minimal disruption to network traffic.

In its default configuration, NSSU upgrades each line card in a switch or Virtual Chassis one at a time. Traffic continues to flow through the other line cards while a line card is being restarted as part of the upgrade. This behavior allows you to minimize disruption to traffic by configuring link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards. When one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.

To reduce the time an NSSU takes, you can configure line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines or on an EX8200 Virtual Chassis.

NOTE: NSSU line-card upgrade groups are not supported for NSSUs on EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, or any mixed Virtual Chassis composed of EX4200, EX4500, or EX4550 switches.

When you define an upgrade group, NSSU upgrades the line cards in the upgrade group at the same time instead of sequentially. 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 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 (NSSU Upgrade Groups)
  [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 (NSSU Upgrade Groups) 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]
  ```

**Related Documentation**

- Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 44
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure) on page 95
- Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) on page 103

- Understanding Nonstop Software Upgrade on EX Series Switches on page 13

## Configuring Power Supply Redundancy (CLI Procedure)

By default, the power management feature in EX6200 and EX8200 switches is configured to manage the power supplies for N+1 redundancy, in which one power supply is held in reserve for backup if any one of the other power supplies is removed or fails.

You can configure power management to manage the power supplies for N+N redundancy. For example, to set up your AC power supplies for dual power feed, N+N redundancy is required. In N+N redundancy, power management allocates half of the online power supplies to normal operating power and half to redundant power. If you have an odd number of online power supplies, power management allocates one more power supply to normal operating power than to redundant power.

This topic describes how to configure power management for N+N redundancy and how to revert back to N+1 redundancy if your deployment needs change.

Before you configure power management for N+N redundancy, ensure that you have sufficient power supplies to meet the power requirements of an N+N configuration. Use the show chassis power-budget-statistics command to display your current power budget.

NOTE: To allow more power to be available to line cards in an EX8200 switch, power management compensates for the reduced normal operating power in an N+N configuration by allocating less power to the chassis than it does in an N+1 configuration. For the EX8208 switch, the power allocated to the chassis is reduced to 1200 W from 1600 W. For the EX8216 switch, it is reduced to 1800 W from 2400 W. In determining whether you have enough power for an N+N configuration, take this reduction of allocated chassis power into account.

The reduction in allocated chassis power is achieved by reducing the maximum fan speed to 60 percent in an N+N configuration from 80 percent in an N+1 configuration. Because the maximum fan speed is reduced, it is possible that a line card that overheats would be shut down sooner in an N+N configuration than in an N+1 configuration.

On EX6200 switches, the same amount of power is allocated for the chassis in N+N configurations as in N+1 configurations.

To configure N+N redundancy:

```
[edit chassis]
user@switch# set psu redundancy n-plus-n
```

To revert back to N+1 redundancy:

```
[edit chassis]
user@switch# delete chassis psu redundancy n-plus-n
```

**Related Documentation**

- Configuring the Power Priority of Line Cards (CLI Procedure) on page 53
- Verifying Power Configuration and Use on page 113
- Understanding Power Management on EX Series Switches on page 21

## Configuring the Power Priority of Line Cards (CLI Procedure)

The power management facility on EX6200 and EX8200 switches allows you to assign power priorities to the slots occupied by line cards. Power management provides power to the slots in priority order, which means that line cards in higher priority slots are more likely to receive power than line cards in lower priority slots if power to the switch is insufficient to power all the line cards.

The power priority you assign to a PoE line card affects both the order in which it receives base power and the order in which it receives PoE power. Base power is allocated first to all line cards in priority order. PoE power is then allocated to the PoE line cards in priority order.

When assigning power priority to slots, keep these points in mind:

- 0 is the highest priority. The number of priority levels depends on the number of slots in a switch—for example, for an EX8208 switch, which has eight slots, you can assign a priority of 0 through 7 to a slot.
- All slots are assigned the lowest priority by default.
- If a group of slots shares the same assigned priority, each slot's power priority within the group is based on its slot number, with the lowest-numbered slots receiving power first. For example, if slot 3 and slot 7 each have an assigned power priority of 2, slot 3 has the higher power priority.
- On EX6200 switches, slots containing a Switch Fabric and Routing Engine (SRE) module are automatically assigned the highest priority. If you assign a priority of 0 to a slot that has a lower number than a slot an SRE module is in, the slot with an SRE module still receives power first. You cannot change the power priority of slot containing an SRE module.

To assign or change the power priority for a slot:

```
[edit chassis]
user@switch# set fpc slot power-budget-priority priority
```

For example, to set slot 6 to priority 0, enter:

```
[edit chassis]
user@switch# set fpc 6 power-budget-priority 0
```

**Related Documentation**

- Configuring Power Supply Redundancy (CLI Procedure) on page 51
- Verifying Power Configuration and Use on page 113
- Understanding Power Management on EX Series Switches on page 21

## Configuring VRRP for IPv6 (CLI Procedure)

By configuring the Virtual Router Redundancy Protocol (VRRP) on EX Series switches, you can enable hosts on a LAN to make use of redundant routing platforms on that LAN

without requiring more than the static configuration of a single default route on the hosts. You can configure VRRP for IPv6 on Gigabit Ethernet, 10-Gigabit Ethernet, and logical interfaces.

To configure VRRP for IPv6:

1. Configure VRRP group support on interfaces:

   ```
   [edit interfaces interface-name unit logical-unit-number family inet6 address address]
   user@switch# set vrrp-inet6-group group-id priority number virtual-inet6-address address virtual-link-local-address ipv6-address
   ```

   You must explicitly define a virtual link local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link local address must be on the same subnet as the physical interface address.

2. If you want to configure the priority order in which this switch functioning as a backup router becomes the master router if the master router becomes nonoperational, configure a priority for this switch:

   ```
   [edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
   user@switch# set priority number
   ```

3. Specify the interval in milliseconds in which the master router sends advertisement packets to the members of the VRRP group:

   ```
   [edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
   user@switch# set inet6-advertise-interval milliseconds
   ```

4. By default, a higher-priority backup router preempts a lower-priority master router.

   - To explicitly enable the master router to be preempted:

     ```
     [edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
     user@switch# set preempt
     ```

   - To prohibit a higher-priority backup router from preempting a lower priority master router:

     ```
     [edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
     user@switch# set no-preempt
     ```

**Related Documentation**

- show vrrp on page 147
- Understanding VRRP on EX Series Switches on page 35

## Determining and Setting Priority for Switches Connected to an EX Series RPS

A Redundant Power System (RPS) provides backup power according to the RPS priority configured on the standalone EX Series switches or Virtual Chassis member switches connected to it. If all switches connected to the RPS are set to the default priority of 1, the priority is determined on the basis of the RPS port to which they are connected, with higher port numbers having the higher priorities.

The number of switches for which an RPS can provide backup power depends on whether the switches provide power over Ethernet (PoE).

- PoE: A fully loaded RPS provides backup power to a maximum of three switches that are enabled for PoE—the result in this case is one switch powered per power supply. If more than three PoE-enabled switches are connected to the RPS and the RPS is already providing backup power to three switches when another switch's power supply fails, the RPS detects this and re-allots backup power as required. It would then stop providing backup power to a low-priority switch to provide backup power to a higher-priority switch.
- Non-PoE: If you changed the RPS power setting to non-PoE with the command request redundant-power-system multi-backup, your RPS is configured to provide back up power to as many as six non-PoE switches on a fully loaded RPS. Each power supply can support two switches when the switches do not need enough power for PoE.

NOTE: Before an RPS can back up a switch connected to it, the switch's RPS status must be ARMED. There are two ways to determine whether a switch's RPS status is ARMED—either check that the corresponding port LED on the RPS is lit and on steady or issue this command from the switch's CLI: show chassis redundant-power-system.

This topic describes how to determine and set the power priority for a switch connected to an RPS.


- Using RPS Default Configuration on page 55
- Setting the EX Series RPS Priority for a Switch (CLI) on page 55


## Using RPS Default Configuration

No configuration is required on an RPS if you:

- Plan to back up as many as six non-PoE switches
- Back up three PoE switches with three RPS power supplies
- Back up four or more PoE switches with RPS three power supplies and let the RPS port to which the switch is connected determine the priority

By default, an RPS assigns priority to switches on the basis of their switch connector port location, with the with higher port numbers having the higher priorities. By default, all switches are themselves configured with the same RPS priority (priority 1, the lowest), which is why priority is derived from the RPS connector port numbers.

## Setting the EX Series RPS Priority for a Switch (CLI)

Each switch connected to RPS has an RPS priority value—that priority value determines which PoE switches receive power first from the RPS. By default, all switches are configured for priority 1 so priority is then determined by switch connector port location, left (lowest) to right (highest).

You can change the priority of a switch to 0 (off), or 1 (lowest) through 6 (highest) from the switch itself—this configuration takes precedence over switch connector port location.

To set or change the priority for a switch that does not support Virtual Chassis:

```
[edit]
user@switch# set redundant-power-system priority
```

To set or change the priority for a switch that supports Virtual Chassis:

```
[edit]
user@switch# set redundant-power-system member vc-member-id priority priority-number
```

Where member is 0 for a switch that has never been configured in a Virtual Chassis.

**Related Documentation**

- Understanding How Power Priority Is Determined and Set for Switches Connected to the EX Series Redundant Power System on page 32
- Connecting the EX Series Redundant Power System to an EX2200 or EX3300 Switch
- EX Series Redundant Power System Hardware Overview on page 29

CHAPTER 10

# Configuration Statements



## [edit chassis] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit chassis]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see EX Series Switch Software Features Overview.

This topic lists:



### Supported Statements in the [edit chassis] Hierarchy Level

The following hierarchy shows the **[edit chassis]** configuration statements supported on EX Series switches:

```
chassis {
    aggregated-devices {
        ethernet {
            device-count number;
        }
    }
    alarm {
        ethernet {
            link-down (ignore | red | yellow);
        }
```

```
    management-ethernet {
      link-down (ignore | red | yellow);
    }
  }
  auto-image-upgrade;
  fpc slot-number{
    pic pic-number {
      sfpplus {
        pic-mode mode;
      tunnel-port port-number tunnel-services;
      }
    }
    power (off | on);
    power-budget-priority priority;
  }
  lcd-menu {
    fpc slot-number {
      menu-item (menu-name | menu-option);
        disable;
  nssu {
    upgrade-group group-name {
      fpcs (slot-number | [list-of-slot-numbers]);
      member member-id {
        fpcs (slot-number | [list-of-slot-numbers]);
      }
    }
  }
  psu {
    redundancy {
      n-plus-n;
    }
  redundancy {
    failover (on-disk-failure | on-loss-of-keepalives);
    graceful-switchover;
    keepalive-time seconds;
  }
  routing-engine {
    on disk-failure
      disk-failure-action action;
    }
  }
  vrf-mtu-check;
}
```

## Unsupported Statements in the [edit chassis] Hierarchy Level

All statements in the **[edit chassis]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 7: Unsupported [edit chassis] Configuration Statements on EX Series Switches**

| Statement | Hierarchy |
|---|---|

NOTE: Variables, such as *filename*, are not shown in the statements or hierarchies.

Table 7: Unsupported [edit chassis] Configuration Statements on EX Series Switches *(continued)*

| Statement | Hierarchy |
|---|---|
| container-devices | [edit chassis] |
| link-down | [edit chassis container-devices] |
| no-multi-rate | [edit chassis fpc pic] |
| pic-mode | [edit chassis fpc pic shdsl] |
| power-budget-priority | [edit chassis fpc] |
| q-pic-large-buffer | edit chassis fpc pic] |
| shdsl | [edit chassis fpc pic] |

**Related Documentation**

- Configuring Aggregated Ethernet Links (CLI Procedure)
- Upgrading Software Using Automatic Software Download
- Configuring the LCD Panel on EX Series Switches (CLI Procedure)
- Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)
- Configuring Power Supply Redundancy (CLI Procedure) on page 51
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 53
- Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure) on page 49

## [edit interfaces] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit interfaces]** hierarchy:

- [edit interfaces ae] Configuration Statement Hierarchy on EX Series Switches
- [edit interfaces ge] Configuration Statement Hierarchy on EX Series Switches
- [edit interfaces interface-range] Configuration Statement Hierarchy on EX Series Switches
- [edit interfaces lo] Configuration Statement Hierarchy on EX Series Switches
- [edit interfaces me] Configuration Statement Hierarchy on EX Series Switches
- [edit interfaces vlan] Configuration Statement Hierarchy on EX Series Switches
- [edit interfaces vme] Configuration Statement Hierarchy on EX Series Switches
- [edit interfaces xe] Configuration Statement Hierarchy on EX Series Switches

**Related Documentation**

- EX Series Switches Interfaces Overview
- Configuring Aggregated Ethernet Links (CLI Procedure)
- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- Configuring a Layer 3 Subinterface (CLI Procedure)
- Configuring Routed VLAN Interfaces (CLI Procedure)
- Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)
- *Junos OS Interfaces Fundamentals Configuration Guide*
- *Junos OS Ethernet Interfaces Configuration Guide*

## [edit protocols vrrp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols vrrp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see EX Series Switch Software Features Overview.

This topic lists:


- Supported Statements in the [edit protocols vrrp] Hierarchy Level on page 60
- Unsupported Statements in the [edit protocols vrrp] Hierarchy Level on page 61


### Supported Statements in the [edit protocols vrrp] Hierarchy Level

The following hierarchy shows the **[edit protocols vrrp]** configuration statements supported on EX Series switches:

```
protocols {
  vrrp {
    failover-delay milliseconds;
    startup-silent-period seconds;
    traceoptions {
      file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
}
```

## Unsupported Statements in the [edit protocols vrrp] Hierarchy Level

All statements in the **[edit protocols vrrp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

**Related Documentation**

- Junos® OS Ethernet Interfaces
- [edit protocols] Configuration Statement Hierarchy on EX Series Switches

## chassis

**Syntax**

```
chassis {
  aggregated-devices {
    ethernet (Aggregated Devices) {
      device-count number;
    }
  }
  auto-image-upgrade;
  fpc slot {
    pic pic-number {
      sfpplus {
        pic-mode  mode;
      }
    }
    power-budget-priority priority;
  }
  lcd-menu {
    fpc slot-number {
      menu-item (menu-name | menu-option) {
        disable;
      }
    }
  }
  nssu {
    upgrade-group group-name {
      fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
      member (NSSU Upgrade Groups) member-id {
        fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
      }
    }
  }
  psu {
    redundancy {
      n-plus-n (Power Management);
    }
  }
  redundancy {
    graceful-switchover;
  }
}
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure chassis-specific properties for the switch.

The remaining statements are explained separately.

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

**Related Documentation**

- Configuring Aggregated Ethernet Links (CLI Procedure)
- Upgrading Software Using Automatic Software Download
- Configuring the LCD Panel on EX Series Switches (CLI Procedure)
- Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)
- Configuring Power Supply Redundancy (CLI Procedure) on page 51
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 53
- Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure) on page 49

## commit synchronize

**Syntax** commit synchronize;

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 10.4 for EX Series switches.

**Description** For devices with multiple Routing Engines only. Configure the **commit** command to automatically result in a **commit synchronize** action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the **commit** command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.

NOTE: When you configure nonstop active routing (NSR), you must include the commit synchronize statement. Otherwise, the commit operation fails.

On the TX Matrix router, synchronization only occurs between the Routing Engines within the same chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the T640 routers. That is, the master Routing Engine on the TX Matrix router distributes the configuration to the master Routing Engine on each T640 router. Likewise, the backup Routing Engine on the TX Matrix router distributes the configuration to the backup Routing Engine on each T640 router.

In EX Series Virtual Chassis configurations:

- On EX4200 switches in Virtual Chassis, synchronization occurs between the switch in the master role and the switch in the backup role.
- On EX8200 switches in a Virtual Chassis, synchronization occurs only between the master and backup XRE200 External Routing Engines.

**Options**

**and-quit**—(Optional) (EX Series only) Quit configuration mode if the commit synchronization succeeds.

**comment**—(Optional) (EX Series only) Write a message to the commit log.

**and-force**—(Optional) (EX Series only) Force a commit synchronization on the other Routing Engine (ignore warnings).

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

**Related Documentation**

- Synchronizing the Routing Engine Configuration

- Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically

## disk-failure-action

**Syntax** disk-failure-action (halt | reboot);

**Hierarchy Level** [edit chassis redundancy on-disk-failure]
[edit chassis routing-engine on-disk-failure]

**Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.

**Description** Configure the Routing Engine to halt or reboot when the Routing Engine hard disk fails.

**Options** **halt**—Specify the Routing Engine to halt.

**reboot**—Specify the Routing Engine to reboot.

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

**Related Documentation**

- graceful-switchover on page 69
- Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors
- Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)
- High Availability Features for EX Series Switches Overview on page 3

## failover (Chassis)

**Syntax**

```
failover {
  on-disk-failure;
  on-loss-of-keepalives;
}
```

**Hierarchy Level** [edit chassis redundancy]

**Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.

**Description** Specify conditions on the master Routing Engine that cause the backup router to take mastership.

The remaining statements are explained separately.

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

**Related Documentation**

- graceful-switchover on page 69
- On Detection of a Hard Disk Error on the Master Routing Engine
- Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)
- High Availability Features for EX Series Switches Overview on page 3

## fpc

**Syntax**

```
fpc slot {
  pic pic-number {
    sfpplus {
      pic-mode mode;
    }
    tunnel-port port-number tunnel-services;
  }
  power (off | on);
  power-budget-priority priority;
}
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description** On an EX3200 or an EX4200 switch, specify the port of the SFP+ uplink module for which you want to configure the operating mode.

On an EX6200 or an EX8200 switch, specify the line card slot for which you want to assign a power priority.

NOTE: On an EX6200 switch, you cannot change the power priority of a slot containing a Switch Fabric and Routing Engine (SRE) module. Although the CLI allows you to set a different power priority for the slot, your change does not go into effect, and the power priority remains 0. A message is sent to the system log to inform you that changing the power priority of the slot is unsupported.

For generic routing encapsulation (GRE) tunneling, use the **fpc** statement with the **tunnel-port** statement to specify the port on the switch that you want to convert to a GRE tunnel port.

**Options** ***slot***—Number of the slot:

- 0—EX3200 and standalone EX4200 switches. The FPC value refers to the switch itself.
- 0–9—EX4200 switch in a Virtual Chassis configuration. The value corresponds to the switch's member ID.
- 0–3 and 6–9—EX6210 switch. The slot is a line card slot.

  4–5—The slot is a line card slot or an SRE module slot.
- 0–7—EX8208 switch. The slot is a line card slot.
- 0–15—EX8216 switch. The slot is a line card slot.

The remaining statements are explained separately.

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

**Related Documentation**
- Setting the Mode on an SFP+ Uplink Module (CLI Procedure)
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 53
- Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)

## fpcs (NSSU Upgrade Groups)

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

**Hierarchy Level**
[edit chassis nssu upgrade-group *group-name*],
[edit chassis nssu upgrade-group *group-name* member (NSSU Upgrade Groups) *member-id*]

**Release Information**
Statement introduced in Junos OS Release 10.4 for EX Series switches.

**Description**
(EX8200 switches and EX8200 Virtual Chassis only) Assign one or more line cards to a line-card upgrade group by specifying their slot numbers.

**Options**
***list-of-slot-numbers***—A list of slot numbers of multiple line cards to be included in the upgrade group. Separate the slot numbers with spaces and enclose the list in square brackets—for example: [3 4 7].

***slot-number***—The slot number of a single line card 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 on page 44
- Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure) on page 49

## graceful-switchover

**Syntax** graceful-switchover;

**Hierarchy Level** [edit chassis redundancy]

**Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.

**Description** For switches with more than one Routing Engine, including those in a Virtual Chassis, configure the master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.

**Default** Graceful Routing Engine switchover (GRES) is disabled.

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

**Related Documentation**

- Example: Configuring Nonstop Active Routing on EX Series Switches on page 41
- Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)
- Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure) on page 47
- Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)

## hold-time

**Syntax** hold-time *seconds*;

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

**Release Information** Statement introduced in Junos OS Release 10.0 for EX Series switches.

**Description** Configure the time in seconds after which a backup router with the highest priority preempts the master router.

**Options** ***seconds***—Hold-time period.

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

**Related Documentation**

- Configuring VRRP for IPv6 (CLI Procedure) on page 53

## inet6-advertise-interval

**Syntax** inet6-advertise-interval *milliseconds*;

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

**Release Information** Statement introduced in Junos OS Release 10.0 for EX Series switches.

**Description** Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.

**Options** ***milliseconds***—Interval, in milliseconds, between advertisement packets.

**Range:** 100 to 40,000 ms

**Default:** 1 second

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

**Related Documentation**

- Configuring VRRP for IPv6 (CLI Procedure) on page 53

 

## keepalive-time

**Syntax** keepalive-time *seconds*;

**Hierarchy Level** [edit chassis redundancy]

**Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.

**Description** Configure the time period that must elapse before the backup router takes mastership when it detects loss of the keepalive signal.

**Default** The **on-loss-of-keepalives** statement at the **[edit chassis redundancy failover]** hierarchy level must be included for failover to occur.

When the **on-loss-of-keepalives** statement is included and graceful Routing Engine switchover *is not* configured, failover occurs after 300 seconds (5 minutes).

When the **on-loss-of-keepalives** statement is included and graceful Routing Engine switchover *is* configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds.

**Options** ***seconds***—Time before the backup router takes mastership when it detects loss of the keepalive signal. The range of values is 2 through 10,000.

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

**Related Documentation**

- failover on page 66
- graceful-switchover on page 69
- on-loss-of-keepalives on page 77
- High Availability Features for EX Series Switches Overview on page 3

## member (NSSU Upgrade Groups)

**Syntax**

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

**Hierarchy Level** [edit chassis nssu upgrade-group *group-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for EX Series switches.

**Description** (EX8200 Virtual Chassis only) Specify the Virtual Chassis member whose line cards you are assigning to the upgrade group.

To create an upgrade group that includes line cards from different Virtual Chassis members, separately configure each Virtual Chassis member as part of the upgrade group.

**Options** *member-id*—The ID of the Virtual Chassis member.

The remaining statement is explained separately.

**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 on page 44
- Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure) on page 49

## member (Redundant Power System)

**Syntax**

```
member vc-member-number {
  priority (0|1|2|3|4|5|6);
}
```

**Hierarchy Level** [edit redundant-power-system]

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

**Description** Specify the Virtual Chassis member ID of a switch connected to the Redundant Power System (RPS) for backup power supply. The member ID is required only for switches that can be configured in a Virtual Chassis. If the switch has never been configured in a Virtual Chassis, the value is always 0.

**Options** *member-number*—Member ID of a switch that has Virtual Chassis capability that is connected to the RPS.

**Range:** 0 through maximum members in the Virtual Chassis

**Default:** 0

The remaining statement is explained separately.

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

**Related Documentation**

- Determining and Setting Priority for Switches Connected to an EX Series RPS on page 54

## n-plus-n (Power Management)

**Syntax**

```
n-plus-n;
```

**Hierarchy Level** [edit chassis psu redundancy]

**Release Information** Statement introduced in Junos OS Release 10.2 for EX Series switches.

**Description** Configure N+N power supply redundancy for power management on an EX6200 or EX8200 switch.

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

**Related Documentation**

- Configuring Power Supply Redundancy (CLI Procedure) on page 51

## nonstop-bridging (Ethernet Switching)

**Syntax** nonstop-bridging;

**Hierarchy Level** [edit ethernet-switching-options]

**Release Information** Statement introduced in Junos OS Release 11.3 for EX Series switches.

**Description** For switches with two Routing Engines or for Virtual Chassis, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 protocol information for the Layer 2 protocols that support nonstop bridging (NSB). For a list of the EX Series switches and Layer 2 protocols that support nonstop bridging, see EX Series Switch Software Features Overview.

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

**Related Documentation**
- Configuring Nonstop Bridging on EX Series Switches (CLI Procedure) on page 48

## nonstop-routing

**Syntax** nonstop-routing;

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

**Release Information** Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description** For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.

**Default** disabled

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

**Related Documentation**
- Configuring Nonstop Active Routing
- Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure) on page 47

## nssu

**Syntax**

```
nssu {
  upgrade-group group-name {
    fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
    member (NSSU Upgrade Groups) member-id {
      fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
    }
  }
}
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 10.4 for EX Series switches.

**Description** (EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis only) Define a line-card upgrade group for nonstop software upgrade (NSSU). All line cards in an upgrade group are upgraded to the new software version at the same time.

The remaining statements are explained separately.

**Default** If no line-card upgrade groups are defined, NSSU upgrades line cards one at a time in ascending order by slot 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 on page 44
- Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure) on page 49

## on-disk-failure

**Syntax**

```
on-disk-failure {
  disk-failure-action (halt | reboot);
}
```

**Hierarchy Level**

[edit chassis redundancy]
[edit chassis routing-engine]

**Release Information**

Statement introduced in Junos OS Release 9.2 for EX Series switches.

**Description**

Instruct the router to halt or reboot if it detects hard disk errors on the Routing Engine.

**Options**

The remaining statement is explained separately.

**Required Privilege Level**

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

**Related Documentation**

- graceful-switchover on page 69
- Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors
- Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)
- High Availability Features for EX Series Switches Overview on page 3

## on-loss-of-keepalives

**Syntax** on-loss-of-keepalives;

**Hierarchy Level** [edit chassis redundancy failover]

**Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.

**Description** Instruct the backup router to take mastership if it detects a loss of keepalive signal from the master Routing Engine.

**Default** The **on-loss-of-keepalives** statement must be included at the **[edit chassis redundancy failover]** hierarchy level for failover to occur.

When the **on-loss-of-keepalives** statement is included but graceful Routing Engine switchover *is not* configured, failover occurs after 300 seconds (5 minutes).

When the **on-loss-of-keepalives** statement is included and graceful Routing Engine switchover *is* configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds.

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

**Related Documentation**

- graceful-switchover on page 69
- keepalive-time on page 71
- Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)
- High Availability Features for EX Series Switches Overview on page 3

## power-budget-priority

**Syntax** power-budget-priority *priority*;

**Hierarchy Level** [edit chassis fpc *slot*]

**Release Information** Statement introduced in Junos OS Release 10.2 for EX Series switches.

**Description** Assign a power priority to the specified line card slot on an EX6200 or EX8200 switch.

NOTE: On an EX6200 switch, you cannot change the power priority of a slot containing a Switch Fabric and Routing Engine (SRE) module. Although the CLI allows you to set a different power priority for the slot, your change does not go into effect, and the power priority remains 0. A message is sent to the system log to inform you that changing the power priority of the slot is unsupported.

**Default** All line card slots are initially assigned the lowest priority, with the exception of slot 4 and slot 5 on the EX6200 switch, which always are assigned a priority of 0.

**Options** ***priority***—Assigned power priority for the slot, with 0 being the highest priority:

- 0 through 9 for an EX6200 switch
- 0 through 7 for an EX8208 switch
- 0 through 15 for an EX8216 switch

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

**Related Documentation**

- Configuring the Power Priority of Line Cards (CLI Procedure) on page 53

 

## preempt

**Syntax**

```
(preempt | no-preempt) {
  hold-time seconds;
}
```

**Hierarchy Level**

[edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

**Release Information**

Statement introduced in Junos OS Release 10.0 for EX Series switches.

**Description**

Configure whether a backup router can preempt a master router:

- **preempt**—Allow the master router to be preempted.
- **no-preempt**—Prohibit the preemption of the master router.

The remaining statement is explained separately.

**Required Privilege Level**

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

**Related Documentation**

- Configuring VRRP for IPv6 (CLI Procedure) on page 53

## priority

**Syntax** priority *number*;

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*],
[edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

**Release Information** Statement introduced in Junos OS Release 10.0 for EX Series switches.

**Description** Configure a switch's priority for becoming the master default routing platform. The routing platform with the highest priority within the group becomes the master.

**Options** *number*—Routing platform's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.

**Range:** 1 through 255

**Default:** 100 (for backup routers)

NOTE: **Priority 255 cannot be assigned to routed VLAN interfaces (RVIs).**

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

**Related Documentation**

- Configuring VRRP for IPv6 (CLI Procedure) on page 53

 

## priority (Redundant Power System)

**Syntax** priority (0|1|2|3|4|5|6);

**Hierarchy Level** [edit redundant-power-system member]
[edit redundant-power-system member *member-number*]

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

**Description** Configure the backup of any switch connected to the Redundant Power System (RPS) using the CLI on each switch. The determines the order in which the RPS supplies backup power to the switches connected to the RPS. 6 is the highest priority and 1 is lowest. Zero means off or no RPS backup.

If the switch is not reconfigured from the CLI, the default priority is 1. In this case, priority is determined by connector location with the rightmost connector having the highest priority.

For switches that can only be used as standalone switches, this hierarchy level is used for configuration:

[edit redundant-power-system]

For switches that can be used either as standalone switches or configured in a Virtual Chassis, this hierarchy level is used for configuration:

[edit redundant-power-system member *vc-member-number*]

If two or more connections are assigned the same , then the power of each connection is determined based on its switch connector port location, with the rightmost port receiving power first.

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

**Related Documentation**

- Determining and Setting Priority for Switches Connected to an EX Series RPS on page 54

 

## psu

**Syntax**

```
psu {
  redundancy {
    n-plus-n (Power Management);
  }
}
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 10.2 for EX Series switches.

**Description** Configure N+N power supply redundancy for power management on an EX6200 or EX8200 switch.

The remaining statements are explained separately.

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

**Related Documentation**

- Configuring Power Supply Redundancy (CLI Procedure) on page 51

 

## redundancy (Graceful Switchover)

**Syntax**

```
redundancy {
  failover {
    on-disk-failure;
    on-loss-of-keepalives;
  }
  graceful-switchover;
}
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.

**Description** Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a standalone EX6200 or EX8200 switch with more than one Routing Engine.

The remaining statements are explained separately.

**Default** Redundancy is enabled for the Routing Engines.

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

**Related Documentation**

- graceful-switchover on page 69
- Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)
- Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)
- High Availability Features for EX Series Switches Overview on page 3

## redundancy (Power Management)

**Syntax**

```
redundancy {
  n-plus-n (Power Management);
}
```

**Hierarchy Level** [edit chassis psu]

**Release Information** Statement introduced in Junos OS Release 10.2 for EX Series switches.

**Description** Configure N+N power supply redundancy for power management on an EX6200 or EX8200 switch.

The remaining statement is explained separately.

**Default** N+1 power supply redundancy is configured by default.

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

**Related Documentation**

- Configuring Power Supply Redundancy (CLI Procedure) on page 51

## redundant-power-system

**Syntax**

EX2200 switch:

```
redundant-power-system {
  priority (0|1|2|3|4|5|6)
  }
}
```

EX3300 switch:

```
redundant-power-system {
  member vc-member-number {
    priority (0|1|2|3|4|5|6)
  }
}
```

**Hierarchy Level** [edit]

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

**Description** Configure Redundant Power System (RPS) member to ensure higher- switches always receive power backup.

The remaining statements are explained separately.

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

**Related Documentation**

- Determining and Setting Priority for Switches Connected to an EX Series RPS on page 54

## routing-engine

**Syntax**

```
routing-engine {
  on-disk-failure {
    disk-failure-action (halt | reboot);
  }
}
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.

**Description** Configure a Routing Engine to halt or reboot automatically when a hard disk error occurs. A hard disk error may cause a Routing Engine to enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. Rebooting or halting prevents this.

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

**Related Documentation**

- graceful-switchover on page 69
- Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors
- High Availability Features for EX Series Switches Overview on page 3
- *Junos High Availability Configuration Guide*

## traceoptions (Routing Options)

**Syntax**

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <disable>;
}
```

**Hierarchy Level**

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options],
[edit routing-options flow],
[edit routing-options multicast]

**Release Information**

Statement introduced before Junos OS Release 7.4.
**nsr-synchronization** flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.
**nsr-synchronization** and **nsr-packet** flags for BFD sessions added in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
**nsr-synchronization** flag for RIP and RIPng added in Junos OS Release 9.0.
**nsr-synchronization** flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.
**nsr-synchronization** flag for PIM added in Junos OS Release 9.3.
**nsr-synchronization** flag for MPLS added in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
**nsr-synchronization** flag for MSDP added in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description**

Define tracing operations that track all routing protocol functionality in the routing device.

To specify more than one tracing operation, include multiple **flag** statements.

**Default**

If you do not include this statement, no global tracing operations are performed.

**Options**

Values:

**disable**—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

**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**. We recommend that you place global routing protocol tracing output in the file **routing-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. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**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 this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files 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 the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

**Required Privilege Level**
routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**
- Example: Tracing Global Routing Protocol Operations
- Tracing Nonstop Active Routing Synchronization Events on page 159

## upgrade-group

**Syntax**

```
upgrade-group group-name {
  fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
  member (NSSU Upgrade Groups) member-id {
    fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
  }
}
```

**Hierarchy Level**
[edit chassis nssu]

**Release Information**
Statement introduced in Junos OS Release 10.4 for EX Series switches.

**Description**
(EX8200 switches and EX8200 Virtual Chassis only) Assign a name to the line-card upgrade group being created for nonstop software upgrade (NSSU).

**Options**
*group-name*—Name of the upgrade group.

The remaining statements are explained separately.

**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 on page 44
- Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure) on page 49

## virtual-inet6-address

**Syntax** virtual-inet6-address [*addresses*];

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

**Release Information** Statement introduced in Junos OS Release 10.0 for EX Series switches.

**Description** Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.

NOTE: The address of an aggregated Ethernet interface (a LAG) or a routed VLAN interface (RVI) cannot be assigned as the virtual router address in a VRRP IPv6 group.

**Options** ***addresses***—Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.

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

**Related Documentation**

- Configuring VRRP for IPv6 (CLI Procedure) on page 53

 

## virtual-link-local-address

**Syntax** virtual-link-local-address *ipv6-address*;

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-inet6-group *group-id*]
[edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

**Release Information** Statement introduced in Junos OS Release 10.0 for EX Series switches.

**Description** Configure a virtual link local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link local address for each VRRP IPv6 group. The virtual link local address must be in the same subnet as the physical interface address.

**Options** ***ipv6-address***—Virtual link local IPv6 address for VRRP for an IPv6 group.

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

**Related Documentation**
- Configuring VRRP for IPv6 (CLI Procedure) on page 53

## vrrp-inet6-group

**Syntax**

```
vrrp-inet6-group group-id {
  inet6-advertise-interval milliseconds;
    preempt{
      hold-time seconds;
    }
    priority number;
    virtual-inet6-address;
    virtual-link-local-address
  }
```

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address*]

**Release Information** Statement introduced in Junos OS Release 10.0 for EX Series switches.

**Description** Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group.

**Options** ***group-id***—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from **00:00:5e:00:01:00** through **00:00:5e:00:01:ff** are reserved for VRRP, as defined in RFC 3768. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

**Range:** 0 through 255

The remaining statements are explained separately.

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

**Related Documentation**

- Configuring VRRP for IPv6 (CLI Procedure) on page 53

 

PART 3

# Administration

CHAPTER 11

# Administration Tasks



## Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)

You can use nonstop software upgrade (NSSU) to upgrade the software on standalone EX6200 or EX8200 switches with redundant Routing Engines. NSSU upgrades the software running on the Routing Engines and line cards with minimal traffic disruption during the upgrade. NSSU is supported on EX8200 switches running Junos OS Release 10.4 or later and on EX6200 switches running Junos OS Release 12.2 or later.

This topic covers:



### Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- (Optional) Configure line-card upgrade groups as described in "Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure)" on page 49. By default, an NSSU upgrades line cards one at a time to allow aggregated Ethernet links that have members on different line cards to remain up through the upgrade process. Configuring line-card upgrade groups reduces the time an upgrade takes because the line cards in each upgrade group are upgraded at the same time rather than sequentially.
- Verify that the Routing Engines are running the same version of the software. Enter the following command:

```
{master}
user@switch> show version invoke-on all-routing-engines
re0:
--------------------------------------------------------------------------
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]

re1:
--------------------------------------------------------------------------
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]
```

If the Routing Engines are not running the same version of the software, use the **request system software add** command to upgrade the Routing Engine that is running the earlier software version. For instructions on upgrading a single Routing Engine, see Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure).

- Ensure that nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

  To verify that nonstop active routing is enabled, execute the following command:

  ```
  {master}
  user@switch> show task replication
          Stateful Replication: Enabled
          RE mode: Master

      Protocol                Synchronization Status
      OSPF                    Complete
      RIP                     Complete
      PIM                     Complete
      RSVP                    Complete
  ```

  If nonstop active routing is not enabled (**Stateful Replication** is **Disabled**), see "Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure)" on page 47 for information on how to enable it.

 

- (Optional) 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 on each Routing Engine to an external storage device with the **request system snapshot** command.

## Upgrading Both Routing Engines Using NSSU

This procedure describes how to upgrade both Routing Engines using NSSU. When the upgrade completes, both Routing Engines are running the new version of the software, and the backup Routing Engine is the new master Routing Engine.

To upgrade both Routing Engines using NSSU:

1. Download the software package by following the procedure in Downloading Software Packages from Juniper Networks.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the **/var/tmp** directory.
3. Log in to the master Routing Engine using the console connection. You can perform an NSSU from the management interface, but a console connection allows you to monitor the progress of the master Routing Engine reboot.
4. Install the new software package:

   ```
   {master}
   user@switch> request system software nonstop-upgrade reboot
    /var/tmp/package-name-m.nZx-distribution.tgz
   ```

   where ***package-name-m.nZx-distribution*.tgz** is, for example, **jinstall-ex-8200-10.4R1.5-domestic-signed.tgz**.

   The switch displays the following status messages as the upgrade executes:

   ```
   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
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


*** FINAL System shutdown message from user@switch ***
System going down IMMEDIATELY

Shutdown NOW!
[pid 2635]
```

NOTE: If you omit the **reboot** option in this step when using an EX8200 switch, you must manually reboot the original master Routing Engine with the **request system reboot** command for the upgrade to complete.

The original master Routing Engine reboots automatically after updating the new master Routing Engine when an NSSU is used to upgrade an EX6200 switch with dual Routing Engines.

5. Log in after the reboot completes. To verify that both Routing Engines have been upgraded, enter the following command:

```
{backup}
user@switch> show version invoke-on all-routing-engines
re0:
--------------------------------------------------------------------------
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]
JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]

re1:
--------------------------------------------------------------------------
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]
JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
```

```
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]
```

6. To verify that the line cards that were online before the upgrade are online after the upgrade, log in to the master Routing Engine and enter the show chassis nonstop-upgrade command:

```
{backup}
user@switch> request routing-engine login master

{master}
user@switch> show chassis nonstop-upgrade
  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)
```

7. If you want to make **re0** the master Routing Engine again, enter the following command:

```
{master}
user@switch> request chassis routing-engine master switch
Toggle mastership between routing engines ? [yes,no] (no) yes
```

   You can verify that **re0** is the master Routing Engine by executing the **show chassis routing-engine** command.

8. To ensure that the resilient dual-root partitions feature operates correctly, execute the following command to copy the new Junos OS image into the alternate root partition on each Routing Engine:

   user@switch> **request system snapshot slice alternate routing-engine both**

   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.

## Upgrading One Routing Engine Using NSSU (EX8200 Switch Only)

This procedure describes how to upgrade one of the Routing Engines using NSSU on an EX8200 switch. When the upgrade completes, the backup Routing Engine is running the new software version and is the new master. The original master Routing Engine, now the backup Routing Engine, continues to run the previous software version.

NOTE: NSSU always upgrades the software on both Routing Engines on an EX6200 switch. Therefore, you cannot upgrade software on one Routing Engine using NSSU on an EX6200 switch.

To upgrade one Routing Engine using NSSU:

1. Download the software package by following the procedure in Downloading Software Packages from Juniper Networks.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the **/var/tmp** directory.
3. Log in to the master Routing Engine.
4. Request an NSSU. On an EX8200 switch, specify the **no-old-master-upgrade** option when requesting the NSSU:

   {master}
   user@switch> **request system software nonstop-upgrade**
   **no-old-master-upgrade /var/tmp/*package-name-m.nZx-distribution*.tgz**

   where ***package-name-m.nZx-distribution*.tgz** is, for example, **jinstall-ex-8200-10.4R2.5-domestic-signed.tgz**.

   The switch displays the following status messages as the upgrade executes:

```
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
```

When the upgrade is complete, the original master Routing Engine (**re0**) becomes the backup Routing Engine.

5. To verify that the original backup Routing Engine (**re1**) has been upgraded, enter the following command:

```
{backup}
user@switch> show version invoke-on all-routing-engines
re0:
--------------------------------------------------------------------------
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]

re1:
--------------------------------------------------------------------------
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]
JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]
```

6. To verify that the line cards that were online before the upgrade are online after the upgrade, log in to the new master Routing Engine and enter the **show chassis nonstop-upgrade** command:

```
{backup}
user@switch> request routing-engine login master

--- JUNOS 12.1-20111229.0 built 2011-12-29 04:12:22 UTC
{master}
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
```

7. To ensure that the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partition of the Routing Engine:

   user@switch> **request system snapshot slice alternate**

   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.

## Upgrading the Original Master Routing Engine (EX8200 Switch Only)

This procedure describes how to upgrade the original master Routing Engine after you have upgraded the original backup Routing Engine as described in "Upgrading One Routing Engine Using NSSU (EX8200 Switch Only)" on page 100 for an EX8200 switch.

1. Log in to the current master Routing Engine (**re1**).
2. Enter configuration mode and disable nonstop active routing:

   ```
   {master}[edit]
   user@switch# delete routing-options nonstop-routing
   ```

3. Deactivate graceful Routing Engine switchover and commit the configuration:

   ```
   {master}[edit]
   user@switch# deactivate chassis redundancy graceful-switchover

   {master}[edit]
   user@switch# commit
   ```

4. Log in to the current backup Routing Engine (**re0**) using a console connection.
5. Request a software installation:

   ```
   user@switch> request system software add reboot
   /var/tmp/package-name-m.nZx-distribution.tgz
   ```

   

   NOTE: When you use NSSU to upgrade only one Routing Engine, the installation package is not automatically deleted from **/var/tmp**, leaving the package available to be used to upgrade the original master Routing Engine.

6. After the upgrade completes, log in to the current master Routing Engine (**re1**) and enter CLI configuration mode.
7. Re-enable nonstop active routing and graceful Routing Engine switchover:

   ```
   [edit]
   user@switch# activate chassis redundancy graceful-switchover
   ```

```
[edit]
user@switch# set routing-options nonstop-routing

[edit]
user@switch# commit
```

8. To ensure that the resilient dual-root partitions feature operates correctly, exit the CLI configuration mode and copy the new Junos OS image into the alternate root partition of the Routing Engine:

   ```
   user@switch> request system snapshot slice alternate
   ```

   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.

9. (Optional) To return control to the original master Routing Engine (**re0**), enter the following command:

   ```
   {master}
   user@switch> request chassis routing-engine master switch
    Toggle mastership between routing engines ? [yes,no] (no) yes
   ```

   You can verify that **re0** is the master Routing Engine by executing the **show chassis routing-engine** command.

**Related Documentation**

- Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 44
- Understanding Resilient Dual-Root Partitions on Switches
- Troubleshooting Software Installation
- Junos OS Package Names
- Understanding Nonstop Software Upgrade on EX Series Switches on page 13
- Understanding Software Installation on EX Series Switches

## Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

You can use nonstop software upgrade (NSSU) to upgrade the software on an EX8200 Virtual Chassis. NSSU upgrades the software running on all Routing Engines with minimal traffic disruption during the upgrade. NSSU is supported on EX8200 Virtual Chassis with redundant XRE200 External Routing Engines running Junos OS Release 11.1 or later.

NOTE: NSSU upgrades all Routing Engines on all members of the Virtual Chassis and on the XRE200 External Routing Engines. Using NSSU, you cannot choose to upgrade the backup Routing Engines only, nor can you choose to upgrade a specific member of the Virtual Chassis. If you need to upgrade a specific member of the Virtual Chassis, see Installing Software for a Single Device in an EX8200 Virtual Chassis.

This topic covers:



## Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- (Optional) Configure line-card upgrade groups as described in "Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure)" on page 49. By default, NSSU upgrades line cards one at a time, starting with the line card in slot 0 of member 0. This permits aggregated Ethernet links that have members on different line cards remain up through the upgrade process. Configuring line-card upgrade groups reduces the time an upgrade takes because the line cards in each upgrade group are upgraded at the same time rather than sequentially.
- Verify that the members are running the same version of the software:

  ```
  {master:8}
  user@external-routing-engine> show version all-members
  ```

  If the 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. For instructions, see Installing Software for a Single Device in an EX8200 Virtual Chassis.

- Ensure that nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

  To verify that nonstop active routing is enabled:

  ```
  {master:8}
  user@switch> show task replication
          Stateful Replication: Enabled
          RE mode: Master

      Protocol                Synchronization Status
          PIM                     Complete
  ```

  If nonstop active routing is not enabled (**Stateful Replication** is **Disabled**), see "Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure)" on page 47 for information on how to enable it.

## Upgrading the Software Using NSSU

This procedure describes how to upgrade the software running on all Routing Engines using NSSU. When the upgrade completes, all Routing Engines are running the new version of the software. The backup external Routing Engine is now the master external Routing Engine, and the internal backup Routing Engines in the member switches are now the internal master Routing Engines in those member switches.

To upgrade all Routing Engines using NSSU:

1. Download the software package for the XRE200 External Routing Engine by following the procedure in Downloading Software Packages from Juniper Networks. The name of the software package for the XRE200 External Routing Engine contains the term **xre200**.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the **/var/tmp** directory.
3. Log in to the master external Routing Engine using the console connection. You can perform an NSSU from the management interface, but a console connection allows you to monitor the progress of the master Routing Engine reboot.
4. Install the new software package:

   ```
   {master:8}
   user@external-routing-engine> request system software nonstop-upgrade reboot
    /var/tmp/package-name-m.nZx-distribution.tgz
   ```

   where ***package-name-m.nZx-distribution*.tgz** is, for example, **jinstall-ex-xre200-11.1R2.5-domestic-signed.tgz**.

   

   NOTE: You can omit reboot option. When you include the reboot option, NSSU automatically reboots the original master Routing Engines after the new image has been installed on them. If you omit the reboot option, you must manually reboot the original master Routing Engines (now the backup Routing Engines) to complete the upgrade. To perform the reboot, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module.

   The switch displays status messages similar to the following messages as the upgrade executes:

   ```
   Chassis ISSU Check Done
   ISSU: Validating Image
   ISSU: Preparing LCC Backup REs
   ISSU: Preparing Backup RE
   Pushing bundle /var/tmp/jinstall-ex-xre200-11.1-20110208.0-domestic-signed.tgz
    to member9
   member9:
   --------------------------------------------------------------------------
   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 member9
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          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 5          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          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 5          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
ISSU: Old Master Upgrade Done
ISSU: IDLE

*** FINAL System shutdown message from root@ ***
System going down IMMEDIATELY


Shutdown NOW!
```

NOTE: If you omit the reboot option in this step, you must complete the upgrade by separately rebooting the original master Routing Engine on each Virtual Chassis member and the original master external Routing Engine. To reboot the original master Routing Engine on a Virtual Chassis member, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module.

5. Log in after the reboot completes. To verify that the software on all Routing Engines in the Virtual Chassis members has been upgraded, enter the following command:

```
{backup:8}
user@external-routing-engine> show version all-members
```

6. Verify that the line cards that were online before the upgrade are online after the upgrade by entering the show chassis nonstop-upgrade command:

```
{backup:8}
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          Online
  FPC 2          Online
  FPC 5          Online
```

**Related Documentation**

- Upgrading Software on an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, or Mixed Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) on page 108
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure) on page 95
- Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 44

- Understanding Resilient Dual-Root Partitions on Switches
- Troubleshooting Software Installation
- Junos OS Package Names
- Understanding Nonstop Software Upgrade on EX Series Switches on page 13
- Understanding Software Installation on EX Series Switches

## Upgrading Software on an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, or Mixed Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

You can use nonstop software upgrade (NSSU) to upgrade the software running on all member switches in most EX Series Virtual Chassis with minimal traffic disruption during the upgrade. NSSU is supported on EX3300 Virtual Chassis, EX4550 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, mixed EX4200, EX4500, and EX4550 Virtual Chassis, and mixed EX4500 and EX4550 Virtual Chassis running Junos OS Release 12.2 or later, and on EX4200 Virtual Chassis, EX4500 Virtual Chassis, or mixed EX4200 and EX4500 Virtual Chassis running Junos OS Release 12.1 or later.

NSSU is also supported on EX8200 Virtual Chassis—see "Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)" on page 103.

This topic covers:



### Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- Ensure that the Virtual Chassis is configured correctly to support NSSU. Verify that:
  - The Virtual Chassis members are connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
  - The 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 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 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 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 nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

  To verify that nonstop active routing is enabled:

```
user@switch> show task replication
        Stateful Replication: Enabled
        RE mode: Master

    Protocol                Synchronization Status
    OSPF                    Complete
    BGP                     Complete
    PIM                     Complete
```

  If nonstop active routing is not enabled (**Stateful Replication** is **Disabled**), see "Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure)" on page 47 for information on how to enable it.

- (Optional) 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 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 backup is the new master.

To upgrade all members using NSSU:

1. Download the software package by following the procedure in Downloading Software Packages from Juniper Networks. If you are upgrading the software running on a mixed Virtual Chassis, download the software packages for both switch types.
2. Copy the software package or packages to the Virtual Chassis. We recommend that you copy the file to the **/var/tmp** directory on the master.
3. Log in to the 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 an EX3300 Virtual Chassis,EX4200 Virtual Chassis, EX4500 Virtual Chassis, or EX4550 Virtual Chassis, enter:

  user@switch> request system software nonstop-upgrade /var/tmp/*package-name*.tgz

  where ***package-name*.tgz** is, for example, **jinstall-ex4200-12.1R2.5-domestic-signed.tgz**.

- On a mixed Virtual Chassis, enter:

  user@switch> **request system software nonstop-upgrade set [/var/tmp/*package-name*.tgz /var/tmp/*package-name*.tgz]**

  where **[/var/tmp/*package-name*.tgz /var/tmp/*package-name*.tgz]** specifies the EX4200 and EX4500 software packages.

The switch displays status messages similar to the following messages as the upgrade executes:

```
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

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)
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 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 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 Resilient Dual-Root Partitions on Switches
- Troubleshooting Software Installation
- Junos OS Package Names
- Understanding Nonstop Software Upgrade on EX Series Switches on page 13
- Understanding Software Installation on EX Series Switches

CHAPTER 12

# Verification Tasks



## Verifying Power Configuration and Use

**Purpose** Verify on an EX6200 or EX8200 switch:

- The power redundancy and line card priority settings
- The PoE power budgets for line cards that support PoE
- Whether the N+1 or N+N power requirements are being met
- Whether the switch has sufficient power for a new line card or an N+N configuration

**Action** Enter the following command:

```
user@switch> show chassis power-budget-statistics
```

Example output for an EX6200 switch:

```
      PSU  0     (EX6200-PWR-AC2500)                :    2500 W   Online
      PSU  1     (EX6200-PWR-AC2500)                :    2500 W   Online
      PSU  2     (EX6200-PWR-AC2500)                :    2500 W   Online
      PSU  3     (EX6200-PWR-AC2500)                :    2500 W   Online
      Total Power supplied by all Online PSUs  :   10000 W
      Power Redundancy Configuration           :     N+1
      Power Reserved for the Chassis           :     500 W
Fan Tray Statistics            Base power    Power Used
FTC  0                         :     300 W      43.04 W
FPC Statistics                  Base power   Power Used   PoE power   Priority

FPC  1   (EX6200-48P)          :     220 W      49.47 W      1440 W       1
FPC  2   (EX6200-48P)          :     220 W      47.20 W       800 W       2
FPC  3   (EX6200-48P)          :     220 W    1493.57 W      1440 W       0
FPC  4   (EX6200-SRE64-4XS)    :     100 W      51.38 W         0 W       0
FPC  5   (EX6200-SRE64-4XS)    :     100 W      50.28 W         0 W       0
FPC  6   (EX6200-48P)          :     220 W      49.38 W       800 W       6
FPC  8   (EX6200-48P)          :     220 W      61.41 W      1440 W       9
FPC  9   (EX6200-48T)          :     150 W      12.49 W         0 W       9

      Total (non-PoE) Power allocated          :    1750 W
      Total Power allocated for PoE            :    5920 W
      Power Available (Redundant case)         :    5750 W
      Total Power Available                    :    2515 W
```

Example output for an EX8200 switch:

```
     PSU  0     (EX8200-AC2K)                  :    1200 W  Online
     PSU  1     (EX8200-AC2K)                  :    1200 W  Online
     PSU  2     (EX8200-AC2K)                  :    1200 W  Online
     PSU  3     (EX8200-AC2K)                  :    1200 W  Online
     Total Power supplied by all Online PSUs   :    4800 W
     Power Redundancy Configuration            :     N+1
     Power Reserved for the Chassis            :    1600 W
FPC Statistics                                  Base power    PoE power   Priority
FPC  0   (EX8200-48T)                         :     350 W          0 W      2
FPC  1   (EX8200-2XS-40P)                     :     387 W        300 W      0
FPC  2   (EX8200-48PL)                        :     267 W        350 W     15
FPC  4   (EX8200-2XS-40P)                     :     387 W        300 W      1
FPC  5   (EX8200-48TL)                        :     230 W          0 W     15
FPC  6   (EX8200-48TL)                        :     230 W          0 W     15

    Total (non-PoE) Power allocated          :    3451 W
    Total Power allocated for PoE            :     950 W
    Power Available (Redundant case)         :     149 W
    Total Power Available                    :     510 W
```

**Meaning**

- Example output for an EX6200 switch —The online power supplies can supply a total of 10,000 W to the switch. The switch is configured for N+1 redundancy, which means 7500 W of redundant power can be supplied. The **Power Available (Redundant case)** field shows that the switch is meeting the N+1 power requirements, with an additional 5750 W available. This value is calculated by subtracting all power allocations except PoE power allocations from redundant power (7500 W).

  The total amount of power available on the switch is 2515 W. This value is calculated by subtracting all power allocations, including PoE power allocations, from the total power (10,000 W). On a switch with PoE line cards, if **Total Power Available** is 0, some or all of the PoE line cards might not be allocated their configured PoE power budgets, which means power to some or all PoE ports might be disabled.

  The power priority order of the line cards, from highest priority line card to the lowest priority line card, is 4, 5, 3, 1, 2, 6, 8, 9. Slots 4 and 5, which contain the Switch Fabric and Routing Engine (SRE) modules, always have highest priority, even if a lower-numbered slot, such as slot 3 in this example, has a priority of 0. Should two or more 2500 W power supplies fail, power management will remove or reduce the PoE power allocations from the PoE line cards in the following order to balance the power budget: 8, 6, 2, 1, and 3.

  The **Power Used** values for the fan tray and line cards shows the actual power being consumed for these components at the time the command was executed. These values are for your information only; power management uses allocated power, which is based on the maximum power the component might consume, and not actual power consumed, in determining its power budget.

- Example output for an EX8200 switch—The online power supplies can supply a total of 4800 W to the switch. The switch is configured for N+1 redundancy, which means 3600 W of redundant power can be supplied. The **Power Available (Redundant case)** field shows that the switch is meeting the N+1 power requirements, with an additional 149 W available. This value is calculated by subtracting all power allocations except

PoE power allocations from redundant power (3600 W). Because 149 W is insufficient power for a line card, another line card cannot be added to the switch while maintaining N+1 redundancy.

The total amount of power available on the switch is 510 W. This value is calculated by subtracting all power allocations, including PoE power allocations, from the total power (4800 W). On a switch with PoE line cards, if **Total Power Available** is 0, some or all of the PoE line cards might not be allocated their configured PoE power budgets, which means power to some or all PoE ports might be disabled.

The power priority order of the line cards, from highest priority line card to the lowest priority line card, is 1, 4, 0, 2, 5, 6. Should one or more 1200 W power supplies fail, power management will remove or reduce the PoE power allocations from the PoE line cards in the following order to balance the power budget: 2, 4, and 1.

**Related Documentation**

- Configuring Power Supply Redundancy (CLI Procedure) on page 51
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 53

CHAPTER 13

# Operational Commands

## request redundant-power-system multi-backup

**Syntax**

EX2200 switch:

```
request redundant-power-system multi-backup
request redundant-power-system no-multi-backup
```

EX3300 switch:

request redundant-power-system multi-backup member *member-number*

request redundant-power-system no-multi-backup member *member-number*

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

**Description** Configure a redundant power system (RPS) to back up six non-Power-over-Ethernet (PoE) powered switches instead of the default which is to back up three PoE-powered switches.

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

**Related Documentation**

- EX Series Redundant Power System Hardware Overview on page 29

**List of Sample Output** request redundant-power-system multi-backup on page 118

### Sample Output

#### request redundant-power-system multi-backup

```
user@switch> request redundant-power-system multi-backup member 1
Sending multi-backup setting to RPS
```

## request system software nonstop-upgrade

**Syntax** request system software nonstop-upgrade (*package-name* | set [*package-name package-name*])
<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.

**Description** Perform a nonstop software upgrade (NSSU) on a switch with redundant Routing Engines or on a Virtual Chassis. The behavior of this command depends on which switch or Virtual Chassis it is executed on:

- When you execute this command on an EX3300, EX4200, EX4500, or EX4550 Virtual Chassis, or a mixed Virtual Chassis composed of any combination of EX4200, EX4500, and EX4550 switches, all Virtual Chassis members are upgraded. The original Virtual Chassis backup becomes the master. The original master is automatically upgraded and rebooted and rejoins the Virtual Chassis 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.

  The original master Routing Engine is automatically rebooted on an EX6200 switch.

  The original master Routing Engine is not automatically rebooted on an EX8200 switch 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 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)" on page 48.

- 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, EX4200, EX4500, EX4550, and mixed Virtual Chassis) or on different line cards (for EX6200 and EX8200 switches, and for EX8200 Virtual Chassis).
- For EX3300, EX4200, EX4500, EX4550, 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 linecard roles are rebooting.
  - The 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 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 linecard 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.

**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.

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

**no-old-master-upgrade**—(Optional) (Standalone 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 switch 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 128

- Upgrading Software on an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, or Mixed Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) on page 108
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure) on page 95
- Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) on page 103

**List of Sample Output**



**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
```

## show chassis nonstop-upgrade

**Syntax** **show chassis nonstop-upgrade**

**Release Information** Command introduced in Junos OS Release 10.4 for EX Series switches.

**Description** (EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis 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 119
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure) on page 95
- Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) on page 103

**List of Sample Output**

show chassis nonstop-upgrade (EX8200 Switch) on page 128
show chassis nonstop-upgrade (EX8200 Virtual Chassis) on page 129

**Output Fields** Table 8 on page 128 lists the output fields for the **show chassis nonstop-upgrade** command. Output fields are listed in the approximate order in which they appear.

Table 8: show chassis nonstop-upgrade Output Fields

| Field Name | Field Description |
|---|---|
| **Item** | Line card slot number. |
| **Status** | State of line card:<br>• **Error**—Line card is in an error state.<br>• **Offline**—Line card is powered down.<br>• **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 power-budget-statistics

**Syntax** **show chassis power-budget-statistics**

**Release Information** Command introduced in Junos OS Release 10.2 for EX Series switches.

**Description** Display the power budget of an EX6200 or an EX8200 switch.

**Required Privilege Level** view

**Related Documentation**

- Verifying Power Configuration and Use on page 113
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 53
- Configuring Power Supply Redundancy (CLI Procedure) on page 51

**List of Sample Output**

show chassis power-budget-statistics (EX6200 Switch) on page 132
show chassis power-budget-statistics (EX8200 Switch) on page 132

**Output Fields** Table 9 on page 130 lists the output fields for the **show chassis power-budget-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 9: show chassis power-budget-statistics Output Fields

| Field Name | Field Description |
|---|---|
| PSU *n* (*supply type*) | Capacity rating of the power supply and whether the power supply is currently operating (**Online**) or not (**Offline**). If a power supply is offline, the capacity is shown as 0 W. |
| Total Power supplied by all Online PSUs | Total number of watts supplied by all currently operating power supplies. |
| Power Redundancy Configuration | Configured power redundancy setting, either N+1 or N+N. |
| Power Reserved for the Chassis | Power reserved for the chassis:<br>• For an EX6200 switch, 500 W.<br>• For an EX8208 switch: 1600 W in an N+1 configuration; 1200 W in an N+N configuration<br>• For an EX8216 switch: 2400 W in an N+1 configuration; 1800 W in an N+N configuration<br>The power reserved for the chassis includes the maximum power requirements for the fan tray and Switch Fabric and Routing Engine (SRE), Routing Engine (RE), and Switch Fabric (SF) modules in both base and redundant configurations. |

Table 9: show chassis power-budget-statistics Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **Fan Tray Statistics** | (EX6200 switch only) Information about the fan tray:<br>• **Base power**—Power allocated to the fan tray in the power budget. This allocation is included in **Power Reserved for the Chassis**.<br>• **Power Used**—Actual power being used by the fan tray. This value is for informational purposes only: the power budget for the switch is based on allocated power (the theoretical maximum the fan tray might use) rather than used power. |
| **FPC *n* (*card type*)** | Information about the line card installed in slot *n*. For EX6200 switches, information about the SRE modules in slot 4 and slot 5 is also shown.<br>• **Base power**—For line cards without PoE ports, the total power allocated to the line card.<br>For line cards with PoE ports, the power allocated to the line card before the PoE power budget is allocated. The base power includes 37 W of PoE power that is always allocated to line cards that support PoE.<br>• **Power Used**—(EX6200 switch only) The actual power being consumed by the line card or SRE module, including PoE power. This value is for informational purposes only: the power budget for the switch is based on allocated power (the theoretical maximum the line card might use) rather than used power.<br>• **PoE power**—For line cards with PoE ports, the PoE power budget allocated to the line card. This value includes the 37 W of PoE power that is always part of the base power allocation for line cards that support PoE.<br>For line cards without PoE ports, the value is always 0 W.<br>• The power priority assigned to the line card slot. |
| **Total (non-PoE) Power allocated** | Power budgeted for all the components in the switch, excluding the PoE power budget allocated to line cards. This value is equal to the power reserved for the chassis plus the base power allocations of all online line cards. |
| **Total Power allocated for PoE** | The total of the PoE power budgets allocated to the line cards in the switch. This figure includes the 37 W of PoE power always included in the base allocation for each line card that supports PoE. |
| **Power Available (Redundant case)** | Unused power available to the switch in the power budget, not including the power reserved for redundancy. If power is insufficient to meet the N+1 or N+N redundancy requirements, this value is 0. PoE power allocations are not included in the calculation of this value. |

Table 9: show chassis power-budget-statistics Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **Total Power Available** | Unused power available to the switch in the power budget. This value is derived by subtracting all power allocations, including PoE power allocations, from the total power available on the switch (the **Total Power supplied by all Online PSUs** value). |

## Sample Output

### show chassis power-budget-statistics (EX6200 Switch)

```
user@switch> show chassis power-budget-statistics
        PSU  0     (EX6200-PWR-AC2500)                 :    2500 W   Online
        PSU  1     (EX6200-PWR-AC2500)                 :    2500 W   Online
        PSU  2     (EX6200-PWR-AC2500)                 :    2500 W   Online
        PSU  3     (EX6200-PWR-AC2500)                 :    2500 W   Online
        Total Power supplied by all Online PSUs  :   10000 W
        Power Redundancy Configuration           :     N+1
        Power Reserved for the Chassis           :     500 W
  Fan Tray Statistics            Base power    Power Used
  FTC  0                         :     300 W      43.04 W
  FPC Statistics                  Base power   Power Used   PoE power   Priority

  FPC  1   (EX6200-48P)          :     220 W      49.47 W      1440 W       1
  FPC  2   (EX6200-48P)          :     220 W      47.20 W       800 W       2
  FPC  3   (EX6200-48P)          :     220 W    1493.57 W      1440 W       0
  FPC  4   (EX6200-SRE64-4XS)    :     100 W      51.38 W         0 W       0
  FPC  5   (EX6200-SRE64-4XS)    :     100 W      50.28 W         0 W       0
  FPC  6   (EX6200-48P)          :     220 W      49.38 W       800 W       6
  FPC  8   (EX6200-48P)          :     220 W      61.41 W      1440 W       9
  FPC  9   (EX6200-48T)          :     150 W      12.49 W         0 W       9

        Total (non-PoE) Power allocated          :    1750 W
        Total Power allocated for PoE            :    5920 W
        Power Available (Redundant case)         :    5750 W
        Total Power Available                    :    2515 W
```

### show chassis power-budget-statistics (EX8200 Switch)

```
user@switch> show chassis power-budget-statistics
    PSU  0     (EX8200-AC2K)                :    2000 W  Online
    PSU  1     (EX8200-AC2K)                :    2000 W  Online
    PSU  2     (EX8200-AC2K)                :    2000 W  Online
    PSU  3     (EX8200-AC2K)                :    2000 W  online
    PSU  4     (EX8200-AC2K)                :    2000 W  Online
    Total Power supplied by all Online PSUs  :   10000 W
    Power Redundancy Configuration           :     N+1
    Power Reserved for the Chassis           :    2400 W
    FPC Statistics                            Base power    PoE power   Priority

    FPC  1   (EX8200-48T)                    :     350 W          0 W     15
    FPC  5   (EX8200-2XS-40P)                :     387 W        792 W      0
    FPC  9   (EX8200-48PL)                   :     267 W        915 W     15
    FPC  10  (EX8200-2XS-40T)                :     350 W          0 W      1
    FPC  12  (EX8200-48T)                    :     350 W          0 W     15
```

```
Total (non-PoE) Power allocated          :    4104 W
Total Power allocated for PoE            :    1707 W
Power Available (Redundant case)         :    3896 W
Total Power Available                    :    4263 W
```

## show chassis redundant-power-system

**Syntax** show chassis redundant-power-system

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

**Description** Display information about the Redundant Power Systems (RPS) connected to the switch.

**Required Privilege Level** view

**Related Documentation**
- Determining and Setting Priority for Switches Connected to an EX Series RPS on page 54

**List of Sample Output**
show chassis redundant-power-system (Standalone Switch) on page 134
show chassis redundant-power-system (Virtual Chassis member) on page 135

**Output Fields** Table 10 on page 134 lists the output fields for the **show chassis redundant-power-system** command. Output fields are listed in the approximate order in which they appear.

### Table 10: show chassis redundant-power-system Output Fields

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Member** | Member number of the switch connected to the RPS—For a switch that has never been configured in a Virtual Chassis, the value is always zero. For a Virtual Chassis member, the range is zero through the maximum number of members in the Virtual Chassis. | All levels |
| **Status** | Status of the RPS:<br>• ARMED—The switch is ready to get backup power from the RPS if power supply fails on the switch.<br>• OFF—The switch has zero and is not configured to receive backup power from the RPS.<br>• BACKED-UP—The switch is receiving power backup from the RPS.<br>• OVER-SUBSCRIBED—The switch cannot receive backup power from the RPS even if you set the . | All levels |
| **RPS** | Serial number of the RPS. | |
| **Port** | Number of the switch connector on the RPS that is connected to a switch. | All levels |

## Sample Output

### show chassis redundant-power-system (Standalone Switch)

```
user@switch> show chassis redundant-power-system

Member Status      RPS            Port
  0    Armed       CG0209121807   0
```

### show chassis redundant-power-system (Virtual Chassis member)

```
user@switch> show chassis redundant-power-system

Member Status      RPS            Port
  0    Armed       CG0209121814   5
  2    Armed       CG0209121815   4
```

## show redundant-power-system led

**Syntax** show redundant-power-system led

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

**Description** Display information about fan status, Redundant Power System (RPS) status, and the switch connectors as displayed by the corresponding LEDs on the RPS.

**Required Privilege Level** view

**Related Documentation**
- LEDs on an EX Series Redundant Power System

**List of Sample Output**
show redundant-power-system led (Standalone Switch) on page 137
show redundant-power-system led (EX3300 Virtual Chassis) on page 137

**Output Fields** Table 11 on page 136 lists the output fields for the **show redundant-power-system led** command. Output fields are listed in the approximate order in which they appear.

### Table 11: show redundant-power-system led Output Fields

| Field Name | Field Description | Level of Output |
|---|---|---|
| **RPS** | The serial number of the RPS. | |
| **RPS Fan** | Status of the RPS power supply fans as displayed by the LED:<br>• Green—All RPS power supply fans are operating fine.<br>• Amber—A fan has failed in at least one RPS power supply. | All levels |
| **RPS System Status** | Status of the RPS system as displayed by the LED:<br>• Green—The RPS is active.<br>• Blinking green—The RPS is booting.<br>• Amber—An RPS power supply has failed.<br>• Off—The RPS is off. | All levels |
| **RPS Port LED Status** | Status of the RPS switch connectors as displayed by the LEDs. These LEDs indicate whether the redundant power source is being used.<br>• Green—The RPS connector is enabled and connected to a switch but the RPS is not actively backing up the switch.<br>• Blinking green—The RPS is backing up the switch connected to the port.<br>• Off—The RPS connector is not connected to a switch.<br>• Amber—The RPS is oversubscribed and the backup power to the switch has failed. | All levels |
| **Port** | Number of one of the six switch connectors on the RPS. | All levels |
| **Status** | Status of each switch connector on the RPS. | All levels |

## Sample Output

### show redundant-power-system led (Standalone Switch)

```
user@switch> show redundant-power-system led

Gathering requested information.
RPS-CG0209121807
    RPS Fan: GREEN
    RPS System Status: GREEN
RPS Port LED Status
Port  Status
   0  GREEN
   1  OFF
   2  OFF
   3  OFF
   4  OFF
   5  OFF
```

### show redundant-power-system led (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system led

Gathering requested information.
RPS-CG0209121814
    RPS Fan: GREEN
    RPS System Status: GREEN
RPS Port LED Status
Port  Status
   0  OFF
   1  OFF
   2  OFF
   3  OFF
   4  OFF
   5  GREEN
RPS-CG0209121815
    RPS Fan: GREEN
    RPS System Status: GREEN
RPS Port LED Status
Port  Status
   0  OFF
   1  OFF
   2  OFF
   3  OFF
   4  GREEN
   5  OFF
```

## show redundant-power-system multi-backup

**Syntax**

show redundant-power-system multi-backup

show redundant-power-system multi-backup member *member-number*

**Release Information**

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

**Description**

Display the current status of the Redundant Power System's (RPS's) ability to back up two switches per power supply when enough power to support Power over Ethernet (PoE) is not needed. This ability is referred to as the RPS's multi-backup ability.

**Required Privilege Level**

view

**Related Documentation**

- request redundant-power-system multi-backup on page 118

**List of Sample Output**

show redundant-power-system multi-backup on page 138

### Sample Output

#### show redundant-power-system multi-backup

```
User@switch> show redundant-power-system multi-backup
Requesting information from redundant-power-system..     Multi-Backup: enabled
```

## show redundant-power-system network

**Syntax** show redundant-power-system network

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

**Description** Display the Redundant Power Supply (RPS) IP address, netmask address, and gateway address required for firmware backup.

**Required Privilege Level** view

**Related Documentation**
- Upgrading Firmware on an EX Series Redundant Power System

**List of Sample Output** show redundant-power-system network on page 139

### Sample Output

#### show redundant-power-system network

```
user@switch> show redundant-power-system network
 Requesting information from redundant-power-system..
    IP Address:   10.93.2.38
    Netmask: 255.255.254.0
    Gateway:  10.93.3.254
```

## show redundant-power-system power-supply

**Syntax** show redundant-power-system power-supply

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

**Description** Display information about the power supplies installed in the Redundant Power System (RPS). After installing a power supply, we recommend that you use this command to be sure that the power supply installed correctly.

**Required Privilege Level** view

**Related Documentation**
- Installing a Power Supply in the EX Series Redundant Power System

**List of Sample Output**
show redundant-power-system power-supply (Standalone Switch) on page 140
show redundant-power-system power-supply (EX3300 Virtual Chassis) on page 140

**Output Fields** Table 12 on page 140 lists the output fields for the **show redundant-power-system power-supply** command. Output fields are listed in the approximate order in which they appear.

Table 12: show redundant-power-system power-supply Output Fields

| Field Name | Field Description | Level of Output |
|---|---|---|
| **RPS** | Serial number of the RPS. | All levels |
| **PSU Slot** | Number of the power supply slot. Slots are numbered 1 through 3. | All levels |
| **Status** | Status of the power supply slots:<br>• Present—The slot contains an RPS power supply.<br>• Empty—The slot is empty. | All levels |
| **Description** | Description of the RPS power supply installed in the slot. | All levels |

### Sample Output

#### show redundant-power-system power-supply (Standalone Switch)

```
user@switch> show redundant-power-system power-supply

Gathering requested information.
RPS-CG0209121807
PSU Slot Status    Description
       1 Online    930W AC
       2 offline     ---
       3 Online    930W AC
```

#### show redundant-power-system power-supply (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system power-supply
```

```
Gathering requested information.
RPS-CG0209121814
PSU Slot Status    Description
       1 Online    930W AC
       2 offline     ---
       3 Online    930W AC
RPS-CG0209121815
PSU Slot Status    Description
       1 Online    930W AC
       2 Online    930W AC
       3 Online    930W AC
```

## show redundant-power-system status

**Syntax** show redundant-power-system status

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

**Description** Display the status information for the switch connectors on the Redundant Power System (RPS).

**Required Privilege Level** view

**Related Documentation**
- Determining and Setting Priority for Switches Connected to an EX Series RPS on page 54
- Installing a Power Supply in the EX Series Redundant Power System

**List of Sample Output**
show redundant-power-system status (Standalone Switch) on page 142
show redundant-power-system status (EX3300 Virtual Chassis) on page 143

**Output Fields** Table 13 on page 142 lists the output fields for the **show redundant-power-system status** command. Output fields are listed in the approximate order in which they appear.

### Table 13: show redundant-power-system status Output Fields

| Field Name | Field Description | Level of Output |
|---|---|---|
| **RPS** | Serial number of the RPS. | All levels |
| **Port** | Number of the switch connector. | All levels |
| **Status** | Status of the switch connector:<br>• ARMED—The switch is ready to get backup power from RPS if power supply fails on the switch.<br>• OFF—The switch has zero and is not configured to receive backup power from RPS.<br>• BACKED-UP—The switch is receiving power backup from RPS.<br>• OVER-SUBSCRIBED—The switch cannot receive backup power from RPS even if you set the . | All levels |
| **Priority** | Priority value of the switch connector. | All levels |
| **Power-Requested** | Power requested by the switch on the corresponding switch connector. | All levels |

## Sample Output

### show redundant-power-system status (Standalone Switch)

```
user@switch> show redundant-power-system status

Gathering requested information.
RPS-CG0209121807
Port Status        Power-requested
```

```
  0  Armed        3         930W
  1  Off          1         ---
  2  Off          1         ---
  3  Off          1         ---
  4  Off          1         ---
  5  Off          1         ---
```

## show redundant-power-system status (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system status

Gathering requested information.
RPS-CG0209121814
Port Status        Power-requested
  0  OFF          1         ---
  1  OFF          1         ---
  2  OFF          1         ---
  3  OFF          1         ---
  4  OFF          1         ---
  5  Armed        5         930W
RPS-CG0209121815
Port Status        Power-requested
  0  OFF          1         ---
  1  OFF          1         ---
  2  OFF          1         ---
  3  OFF          1         ---
  4  Armed        4         930W
  5  OFF          1         ---
```

## show redundant-power-system upgrade

**Syntax** show redundant-power-system upgrade

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

**Description** Display RPS firmware upgrade status (pass or fail), previous RPS firmware version, and current RPS firmware version.

**Required Privilege Level** view

**Related Documentation**
- request redundant-power-system multi-backup on page 118

**List of Sample Output** show redundant-power-system upgrade on page 144

**Output Fields** Table 14 on page 144 lists the output fields for the **show redundant-power-system status** command. Output fields are listed in the approximate order in which they appear.

Table 14: show redundant-power-system upgrade Output Fields

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Firmware Upgrade Status** | Indicates whether the upgrade passed or failed | All levels |
| **Previous Firmware Version** | Firmware version before the upgrade | All levels |
| **Current Firmware Version** | Firmware version after the upgrade | |

### Sample Output

#### show redundant-power-system upgrade

```
user@switch> show redundant-power-system upgrade
Requesting information from redundant-power-system..
Firmware Upgrade Status:   Pass
Previous Firmware Version: 1.0
Current Firmware Version:  1.0
```

## show redundant-power-system version

**Syntax** show redundant-power-system version

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

**Description** Display version information about the Redundant Power System (RPS).

**Required Privilege Level** view

**Related Documentation**
- Installing a Power Supply in the EX Series Redundant Power System
- Packing an EX Series Redundant Power System or Redundant Power System Components for Shipping

**List of Sample Output**
show redundant-power-system version (Standalone Switch) on page 145
show redundant-power-system version (EX3300 Virtual Chassis) on page 145

**Output Fields** Table 15 on page 145 lists the output fields for the **show redundant-power-system version** command. Output fields are listed in the approximate order in which they appear.

### Table 15: show redundant-power-system version Output Fields

| Field Name | Field Description | Level of Output |
|---|---|---|
| **RPS** | Serial number of the RPS. | All levels |
| **Model** | Model name of the RPS. | All levels |
| **RPS Firmware Version** | Version number of the firmware installed on the RPS. | All levels |
| **RPS U-Boot Version** | Version of the bootup software installed on the RPS. | All levels |

## Sample Output

### show redundant-power-system version (Standalone Switch)

```
user@switch> show redundant-power-system version

   RPS-CG0209121807
     Model: EX-PWR_RPS200
     RPS Firmware Version [1.0]
     RPS U-Boot   Version [1.1.6]
```

### show redundant-power-system version (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system version

   RPS-CG0209121814
     Model: EX-PWR_RPS200
     RPS Firmware Version [1.0]
     RPS U-Boot   Version [1.1.6]
   RPS-CG0209121815
```

```
Model: EX-PWR_RPS200
RPS Firmware Version [1.0]
RPS U-Boot   Version [1.1.6]
```

## show vrrp

**Syntax**
show vrrp
<brief | detail | extensive | summary>
<interface *interface-name*>
<track interfaces>

**Release Information**
Statement introduced in Junos OS Release 10.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description**
Display information and status about VRRP groups.

**Options**
**none**—(Same as brief) Display brief status information about all VRRP interfaces.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**interface *interface-name*** —(Optional) Display information and status about the specified VRRP interface.

**track interfaces**—(Optional) Display information and status about VRRP track interfaces.

**Required Privilege Level**
view

**Related Documentation**
- Configuring VRRP for IPv6 (CLI Procedure) on page 53

**List of Sample Output**


**Output Fields**
Table 16 on page 147 lists the output fields for the **show vrrp** command. Output fields are listed in the approximate order in which they appear.

### Table 16: show vrrp Output Fields

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Interface** | Name of the logical interface. | **none, brief, extensive, summary** |
| **Interface index** | Physical interface index number, which reflects its initialization sequence. | **extensive** |
| **Groups** | Total number of VRRP groups configured on the interface. | **extensive** |

Table 16: show vrrp Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Active** | Total number of VRRP groups that are active (that is, whose interface state is either up or down). | **extensive** |
| **Interface VRRP PDU statistics** | Nonerrored statistics for the logical interface:<br>• **Advertisement sent**—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted.<br>• **Advertisement received**—Number of VRRP advertisement PDUs received by the interface.<br>• **Packets received**—Number of VRRP packets received for VRRP groups on the interface.<br>• **No group match received**—Number of VRRP packets received for VRRP groups that do not exist on the interface. | **extensive** |
| **Interface VRRP PDU error statistics** | Errored statistics for the logical interface:<br>• **Invalid IPAH next type received**—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets.<br>• **Invalid VRRP ttl value received**—Number of packets received whose IP time-to-live (TTL) value is not 255.<br>• **Invalid VRRP version received**—Number of packets received whose VRRP version is not 2.<br>• **Invalid VRRP pdu type received**—Number of packets received whose VRRP PDU type is not 1.<br>• **Invalid VRRP authentication type received**—Number of packets received whose VRRP authentication is not none, simple, or md5.<br>• **Invalid VRRP IP count received**—Number of packets received whose VRRP IP count exceeds 8.<br>• **Invalid VRRP checksum received**—Number of packets received whose VRRP checksum does not match the calculated value. | **extensive** |
| **Physical interface** | Name of the physical interface. | **detail, extensive** |
| **Unit** | Logical unit number. | All levels |
| **Address** | Address of the physical interface. | **none, brief, detail, extensive** |
| **Index** | Physical interface index number, which reflects its initialization sequence. | **detail, extensive** |
| **SNMP ifIndex** | SNMP index number for the physical interface. | **detail, extensive** |
| **VRRP-Traps** | Status of VRRP traps: **Enabled** or **Disabled**. | **detail, extensive** |

Table 16: show vrrp Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Type and Address** | Identifier for the address and the address itself:<br>• **lcl**—Configured local interface address.<br>• **mas**—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router.<br>• **vip**—Configured virtual IP addresses. | **none, brief, summary** |
| **Interface state** or **Int state** | State of the physical interface:<br>• **down**—The device is present and the link is unavailable.<br>• **not present**—The interface is configured, but no physical device is present.<br>• **unknown**—The VRRP process has not had time to query the kernel about the state of the interface.<br>• **up**—The device is present and the link is established. | **none, brief, extensive, summary** |
| **Group** | VRRP group number. | **none, brief, extensive, summary** |
| **State** | VRRP state:<br>• **backup**—The interface is acting as the backup router interface.<br>• **bringup**—VRRP is just starting, and the physical device is not yet present.<br>• **idle**—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.<br>• **initializing**—VRRP is initializing.<br>• **master**—The interface is acting as the master router interface.<br>• **transition**—The interface is changing between being the backup and being the master router. | **extensive** |
| **Priority** | Configured VRRP priority for the interface. | **detail, extensive** |
| **Advertisement interval** | Configured VRRP advertisement interval. | **detail, extensive** |
| **Authentication type** | Configured VRRP authentication type: **none**, **simple**, or **md5**. | **detail, extensive** |
| **Preempt** | Whether preemption is allowed on the interface: **yes** or **no**. | **detail, extensive** |
| **Accept-data mode** | Whether the interface is configured to accept packets destined for the virtual IP address: **yes** or **no**. | **detail, extensive** |
| **VIP count** | Number of virtual IP addresses that have been configured on the interface. | **detail, extensive** |
| **VIP** | List of virtual IP addresses configured on the interface. | **detail, extensive** |
| **Advertisement timer** | Time until the advertisement timer expires. | **detail, extensive** |

Table 16: show vrrp Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Master router** | IP address of the interface that is acting as the master. If the VRRP interface is down, the output is **N/A**. | **detail, extensive** |
| **Virtual router uptime** | Time that the virtual router has been up. | **detail, extensive** |
| **Master router uptime** | Time that the master router has been up. | **detail, extensive** |
| **Virtual MAC** | MAC address associated with the virtual IP address. | **detail, extensive** |
| **Tracking** | Whether tracking is **enabled** or **disabled**. | **detail, extensive** |
| **Current priority** | Current operational priority for being the VRRP master. | **detail, extensive** |
| **Configured priority** | Configured base priority for being the VRRP master. | **detail, extensive** |
| **Priority hold-time** | Minimum time interval, in seconds, between successive changes to the current priority. **Disabled** indicates no minimum interval. | **detail, extensive** |
| **Remaining-time** | (**track** option only) Displays the time remaining in the priority hold-time interval. | **detail** |
| **Interface tracking** | Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces. | **detail extensive** |
| **Interface/Tracked interface** | Name of the tracked interface. | **detail extensive** |
| **Int state/Interface state** | Current operational state of the tracked interface: **up** or **down**. | **detail, extensive** |
| **Int speed/Speed** | Current operational speed, in bits per second, of the tracked interface. | **detail, extensive** |
| **Incurred priority cost** | Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority. | **detail, extensive** |
| **Threshold** | Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred.<br><br>An entry of **down** means that the corresponding priority cost is incurred when the interface is down. | **detail, extensive** |
| **Route tracking** | Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes. | **detail, extensive** |
| **Route count** | The number of routes being tracked. | **detail, extensive** |

Table 16: show vrrp Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Route** | The IP address of the route being tracked. | **detail, extensive** |
| **VRF name** | The VPN routing and forwarding (VRF) routing instance that the tracked route is in. | **detail, extensive** |
| **Route state** | The state of the route being tracked: **up**, **down**, or **unknown**. | **detail, extensive** |
| **Priority cost** | Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down. | **detail, extensive** |
| **Active** | Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred. | **detail, extensive** |
| **Group VRRP PDU statistics** | Number of VRRP advertisements sent and received by the group. | **extensive** |
| **Group VRRP PDU error statistics** | Errored statistics for the VRRP group:<br>• **Bad authentication type received**—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be **none**, **simple**, or **md5** and must be the same for all routers in the VRRP group.<br>• **Bad password received**—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group<br>• **Bad MD5 digest received**—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router.<br>• **Bad advertisement timer received**—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group.<br>• **Bad VIP count received**—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance.<br>• **Bad VIPADDR received**—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance. | **extensive** |
| **Group state transition statistics** | State transition statistics for the VRRP group:<br>• **Idle to master transitions**—Number of times that the VRRP instance transitioned from the idle state to the master state.<br>• **Idle to backup transitions**—Number of times that the VRRP instance transitioned from the idle state to the backup state.<br>• **Backup to master transitions**—Number of times that the VRRP instance transitioned from the backup state to the master state.<br>• **Master to backup transitions**—Number of times that the VRRP instance transitioned from the master state to the backup state. | **extensive** |
| **Vlan-id** | ID of Vlan | **detail** |

Table 16: show vrrp Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **VR state** | VRRP information:<br>• **backup**—The interface is acting as the backup router interface.<br>• **bringup**—VRRP is just starting, and the physical device is not yet present.<br>• **idle**—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.<br>• **initializing**—VRRP is initializing.<br>• **master**—The interface is acting as the master router interface.<br>• **transition**—The interface is changing between being the backup and being the master router. | **none, brief** |
| **Timer** | VRRP timer information:<br>• **A**—Time, in seconds, until the advertisement timer expires.<br>• **D**—Time, in seconds, until the Master is Dead timer expires. | **none, brief** |

## Sample Output

### show vrrp

```
user@host> show vrrp
Interface         State        Group    VR state    Timer   Type  Address
ge-0/0/0.121      up             1     master      A 1.052  lcl   gec0::12:1:1:1

                                                             vip   ge80::12:1:1:99

                                                             vip   gec0::12:1:1:99
ge-0/0/2.131      up               1   master      A 0.364  lcl   gec0::13:1:1:1

                                                             vip   ge80::13:1:1:99

                                                             vip   gec0::13:1:1:99
```

### show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see show vrrp on page 152.

### show vrrp detail (IPv6)

```
user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
  Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
```

```
Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

  Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
  Advertisement timer: 0.327s, Master router: ge80::13:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
```

## show vrrp detail (Route Track)

```
user@host> show vrrp detail
Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24
  Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled
  Interface state: up, Group: 100, State: master
  Priority: 150, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
  Advertisement timer: 1.218s, Master router: 30.30.30.30
  Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
  Virtual MAC: 00:00:5e:00:01:64
  Tracking: enabled
    Current priority: 150, Configured priority: 150
    Priority hold-time: disabled
    Interface tracking: disabled
    Route tracking: enabled, Route count: 1
      Route               VRF name       Route state   Priority cost
      192.168.40.0/22     default        up            30
```

## show vrrp extensive

```
user@host> show vrrp extensive
Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          188
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP TTL value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP PDU type received          :            0
    Invalid VRRP authentication type received:           0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
  Advertisement timer: 1.034s, Master router: ge80::12:1:1:1
  Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
```

```
    Advertisement sent              :          188
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
    Idle to master transitions      :            0
    Idle to backup transitions      :            1
    Backup to master transitions    :            1
    Master to backup transitions    :            0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                       :          186
    Advertisement received                   :            0
    Packets received                         :            0
    No group match received                  :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received          :            0
    Invalid VRRP TTL value received          :            0
    Invalid VRRP version received            :            0
    Invalid VRRP PDU type received           :            0
    Invalid VRRP authentication type received:            0
    Invalid VRRP IP count received           :            0
    Invalid VRRP checksum received           :            0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

  Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
  Advertisement timer: 0.396s, Master router: ge80::13:1:1:1
  Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent              :          186
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
    Idle to master transitions      :            0
    Idle to backup transitions      :            1
    Backup to master transitions    :            1
  Master to backup transitions    :            0
```

## show vrrp interface

```
user@host> show vrrp interface
```

```
Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          205
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP TTL value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP PDU type received          :            0
    Invalid VRRP authentication type received:           0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
  Advertisement timer: 0.789s, Master router: ge80::12:1:1:1
  Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent              :          205
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
    Idle to master transitions      :            0
    Idle to backup transitions      :            1
    Backup to master transitions    :            1
    Master to backup transitions    :            0
```

## show vrrp summary

```
user@host> show vrrp summary
Interface     State       Group   VR state      Type   Address
ge-4/1/0.0    up              1   backup        lcl    10.57.0.2
                                                vip    10.57.0.100
```

## show vrrp track detail

```
user@host> show vrrp track detail
Tracked interface: ae1.211
  State: up, Speed: 400m
  Incurred priority cost: 0
  Threshold   Priority cost   Active
  400m        10
  300m        60
  200m        110
  100m        160
  down        190
```

```
  Tracking VRRP interface: ae0.210, Group: 1
    VR State: master
    Current priority: 200, Configured priority: 200
    Priority hold-time: disabled,    Remaining-time: 50.351
```

### show vrrp track summary

```
user@host> show vrrp track summary
Track if      State     Speed       VRRP if   Group   VR State   Current priority
ae1.211       up        400m        ae0.210     1      master     200
```

PART 4

# Troubleshooting

CHAPTER 14

# Troubleshooting Procedures



## Tracing Nonstop Active Routing Synchronization Events

To track the progress of nonstop active routing synchronization between Routing Engines, you can configure nonstop active routing trace options flags for each supported protocol and for BFD sessions and record these operations to a log file.

To configure nonstop active routing trace options for supported routing protocols, include the **nsr-synchronization** statement at the **[edit protocols** *protocol-name* **traceoptions flag]** hierarchy level and optionally specify one or more of the **detail**, **disable**, **receive**, and **send** options:

```
[edit protocols]
bgp {
    traceoptions {
        flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
}
isis {
    traceoptions {
        flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
}
ldp {
    traceoptions {
        flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
}
mpls {
    traceoptions {
        flag nsr-synchronization;
        flag nsr-synchronization-detail;
    }
}
msdp {
    traceoptions {
        flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
```

```
}
(ospf | ospf3) {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
rip {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
ripng {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
pim {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
```

To configure nonstop active routing trace options for BFD sessions, include the **nsr-synchronization** and **nsr-packet** statements at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```
[edit protocols]
bfd {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-packet;
  }
}
```

To trace the Layer 2 VPN signaling state replicated from routes advertised by BGP, include the **nsr-synchronization** statement at the **[edit routing-options traceoptions flag]** hierarchy level. This flag also traces the label and logical interface association that VPLS receives from the kernel replication state.

```
[edit routing-options]
traceoptions {
  flag nsr-synchronization;
}
```

**Related Documentation**

- Configuring Nonstop Active Routing
- Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure) on page 47
- Example: Configuring Nonstop Active Routing
- Example: Configuring Nonstop Active Routing on EX Series Switches on page 41

## Troubleshooting the EX Series Redundant Power System Power On and Power Backup Issues

This topic provides troubleshooting information for problems related to the EX Series Redundant Power System (RPS). This topic describes:



### The EX Series RPS Is Not Powering On

**Problem** The RPS does not power on even though it has a power supply installed and is connected to an AC power source outlet.

The RPS with one EX-PWR3-930-AC power supply installed in it is connected to a switch.

The SYS LED on the power supply side of the RPS is off, and when you check the RPS status using the CLI command show chassis redundant-power-system, the message **No RPS connected** is displayed.

**Cause** A power supply must be installed in the middle slot on the RPS to power on the RPS.

**Solution** Install a power supply in the middle slot on the power supply side of the RPS and verify that the AC power source outlet is properly connected to it. See Installing a Power Supply in the EX Series Redundant Power System.

Verify that the **AC OK** LED and the **DC OK** LED on the power supply in the RPS are lit green.

### A Switch Is Not Recognized by the RPS

**Problem** I can't set up the RPS.

**Cause** A switch must be active to be recognized by the RPS.

**Solution** Activate the switch by configuring it and issuing a commit statement.

### An Error Message Indicates That an RPS Power Supply is Not Supported

**Problem** An RPS error message indicates that an RPS power supply is not supported.

**Cause** RPS supports only one power supply, the EX-PWR3-930-AC. If you install another similar power supply, it may fit in the slot but it is not compatible with RPS.

**Solution** The power supply shipped with your RPS (in a separate box) is an EX-PWR3-930-AC. If you installed more power supplies, you ordered them separately. Replace any other power supply model (such as the EX-PWR2-930-AC) with an EX-PWR3-930-AC model.

## The EX Series Redundant Power System Is Not Providing Power Backup to a Connected Switch

**Problem** The RPS does not provide power backup to a connected switch.

The RPS has an EX-PWR3-930-AC power supply installed in the middle power supply slot and is connected to two switches with power loss, one connected to RPS switch connector port 1 and the other on port 2.

The status LED on the associated switch connector port is not blinking green—it is either solid green (connected) or not lit (off).

**Cause** The RPS provides backup power based on the power priority assigned to each switch.

**Solution** If the status LED on a switch connector port is off, ensure that the RPS cable is properly connected to both the RPS and the switch, and ensure that the priority configured for the switch is not 0. See **show redundant-power-system status**.

If the status LED on switch connector port 1 is on and is steadily green, check the backup priority configured for the switch and assign it a higher priority. See "Determining and Setting Priority for Switches Connected to an EX Series RPS" on page 54

If the status LED on switch connector port 1 is amber, check if the RPS has enough power supplies installed in it to provide backup power. if it does not, install a power supply in an empty power supply slot on the RPS. See Installing a Power Supply in the EX Series Redundant Power System.

If the status LED on switch connector port 1 is still off, check the priority configured for the switch. Ensure that the is not set to 0, which means off. See **show redundant-power-system status**. The priority assigned must be from 1 through 6. See "Determining and Setting Priority for Switches Connected to an EX Series RPS" on page 54.

Verify that a dedicated power supply is installed in the switch. The RPS cannot boot a switch that does not have a dedicated power supply. See Installing a Power Supply in the EX Series Redundant Power System.

Also keep in mind that when the command **request redundant-power-system multi-backup** has been set, support for switches that supply PoE is not guaranteed. To reverse this setting, use the command **request redundant-power-system no-multi-backup**.

## The Wrong Switches Are Being Backed Up

**Problem** Four or more switches are connected to an RPS with three power supplies. When all four switches fail, the wrong three switches have .

 

Four or more switches are connected to an RPS with three power supplies. One or more switches provide PoE to other devices.

When all four switches fail, the wrong three switches have .

**Cause** The RPS provides backup power based on the power priority assigned to each switch. This is derived from two configurations, one of which has precedence over the other one. Initial is derived from the location of the port used to attach a switch—the leftmost connector has lowest priority and the rightmost connector has highest priority. The second, dominant priority configuration is derived from a CLI priority setting on the switch itself. With this CLI configuration, 6 is highest priority and 1 is the lowest priority.

**Solution** Connect the three switches to the three rightmost connectors on the RPS. Then, using the CLI on each switch, set each switch's priority to 1 using the redundant-power-system configuration command **redundant-power-system 1**. Now, physical connection location is determining .

If you do not want to change the cabling on the switches, you can use the configuration statement **redundant-power-system** on all four switches, assigning priority **6** (highest), **5**, **4** and **3** to the appropriate switches. Priority configuration on the switch always overcomes set by connector location.

## Six Switches That Do Not Require PoE Are Not All Being Backed Up

**Problem** Only three switches out of six are simultaneously backed up when all switches experience power supply failure. None of these switches supply PoE power to any device.

The RPS with three EX-PWR3-930-AC power supplies installed in it is connected to six switches, none of which is connected to a non-PoE device.

Only three switches out of six are simultaneously backed up when all switches experience power supply failure. None of these switches supply PoE power to any device.

**Cause** Each power supply can support two switches that do not need enough power for PoE, as long as you configure the RPS to do so.

**Solution** From any of the attached switches, issue the request redundant-power-system multi-backup command from the operational mode. Now standard power will be supplied to two non-PoE switches per power supply.

**Related Documentation**

- Installing a Power Supply in the EX Series Redundant Power System
- Determining and Setting Priority for Switches Connected to an EX Series RPS on page 54
- LEDs on an EX Series Redundant Power System
- EX Series Redundant Power System Hardware Overview on page 29