



High Availability Feature Guide for EX2300, EX3400, and EX4300 Switches

Release
15.1



Modified: 2016-06-30

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

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

High Availability Feature Guide for EX2300, EX3400, and EX4300 Switches
Release 15.1
Copyright © 2016, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <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

	About the Documentation	xi
	Documentation and Release Notes	xi
	Supported Platforms	xi
	Using the Examples in This Manual	xi
	Merging a Full Example	xii
	Merging a Snippet	xii
	Documentation Conventions	xiii
	Documentation Feedback	xv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xv
	Opening a Case with JTAC	xvi
Part 1	Overview	
Chapter 1	High Availability Features Overview	3
	High Availability Features for EX4300 Switches Overview	3
	Redundant Routing Engines	3
	Virtual Chassis	4
	VRRP	4
	Graceful Routing Engine Switchover	4
	Link Aggregation	5
	Nonstop Active Routing and Nonstop Bridging	5
Chapter 2	Nonstop Active Routing Overview	7
	Understanding Nonstop Active Routing on EX Series Switches	7
Chapter 3	Nonstop Bridging Overview	9
	Understanding Nonstop Bridging on EX Series Switches	9
Chapter 4	Nonstop Software Upgrade Overview	11
	Understanding Nonstop Software Upgrade on EX Series Switches	11
	Requirements for Performing an NSSU	12
	How an NSSU Works	14
	EX3300, EX4200, EX4300, EX4500, and Mixed Virtual Chassis	14
	EX6200 and EX8200 Switches	14
	EX8200 Virtual Chassis	16
	NSSU Limitations	17
	NSSU and Junos OS Release Support	17
	Overview of NSSU Configuration and Operation	18

Chapter 5	Power Management Overview	19
	Understanding Power Management on EX Series Switches	19
	Power Priority of Line Cards	20
	How a Line Card's Power Priority Is Determined	20
	Line Card Priority and Line Card Power	21
	Line Card Priority and PoE Power	21
	Line Card Priority and Changes in the Power Budget	22
	Power Supply Redundancy	23
Chapter 6	VRRP Overview	27
	Understanding VRRP on EX Series Switches	27
	Overview of VRRP on EX Series Switches	27
	Examples of VRRP Topologies	28
Part 2	Configuration	
Chapter 7	Configuration Tasks	33
	Configuring Nonstop Active Routing on Switches	33
	Configuring Nonstop Bridging on Switches (CLI Procedure)	35
	Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)	36
	Preparing the Switch for Software Installation	37
	Upgrading the Software Using NSSU	38
	Configuring Power Supply Redundancy (CLI Procedure)	40
	Configuring VRRP for IPv6 (CLI Procedure)	41
Chapter 8	Configuration Examples	43
	Example: Configuring Nonstop Active Routing on Switches	43
Chapter 9	Configuration Statements	47
	[edit chassis] Configuration Statement Hierarchy on EX Series Switches	48
	Supported Statements in the [edit chassis] Hierarchy Level	48
	[edit interfaces] Configuration Statement Hierarchy on EX Series Switches	49
	[edit protocols layer2-control] Hierarchy Level	50
	[edit protocols vrrp] Hierarchy Level	51
	chassis	52
	synchronize	54
	failover (Chassis)	56
	fpc	57
	graceful-switchover	58
	hold-time	59
	inet6-advertise-interval	59
	keepalive-time	60
	n-plus-n (Power Management)	61
	nonstop-bridging	61
	nonstop-routing	62
	preempt	63
	priority	64
	psu	65

	redundancy (Graceful Switchover)	66
	redundancy (Power Management)	67
	traceoptions	68
	vcp-no-hold-time	71
	virtual-inet6-address	72
	virtual-link-local-address	73
	vrrp-inet6-group	74
Part 3	Administration	
Chapter 10	Verification Tasks	77
	Verifying Power Configuration and Use	77
Chapter 11	Operational Commands	81
	request system software nonstop-upgrade	82
	show chassis power-budget-statistics	92
	show vrrp	96
Chapter 12	Troubleshooting Procedures	107
	Tracing Nonstop Active Routing Synchronization Events	107

List of Figures

Part 1	Overview	
Chapter 6	VRRP Overview	27
	Figure 1: Basic VRRP on EX Series Switches	28
	Figure 2: VRRP on Virtual Chassis Switches	29

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiii
Part 1	Overview	
Chapter 4	Nonstop Software Upgrade Overview	11
	Table 3: Platform and Release Support for NSSU	17
Chapter 5	Power Management Overview	19
	Table 4: Available Operating Power in N+1 and N+N Redundancy Configurations	24
Part 3	Administration	
Chapter 11	Operational Commands	81
	Table 5: show chassis power-budget-statistics Output Fields	92
	Table 6: show vrrp Output Fields	96

About the Documentation

- Documentation and Release Notes on page xi
- Supported Platforms on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

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.xml;
    }
  }
}
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.xml; }
```

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.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

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: user@host> configure

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles.	<ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(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 (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <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:
<http://kb.juniper.net/InfoCenter/>
- 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

- [High Availability Features Overview on page 3](#)
- [Nonstop Active Routing Overview on page 7](#)
- [Nonstop Bridging Overview on page 9](#)
- [Nonstop Software Upgrade Overview on page 11](#)
- [Power Management Overview on page 19](#)
- [VRRP Overview on page 27](#)

CHAPTER 1

High Availability Features Overview

- [High Availability Features for EX4300 Switches Overview on page 3](#)

High Availability Features for EX4300 Switches Overview

High availability refers to the hardware and software components that provide redundancy and reliability for network communications. This topic covers the following high availability features of Juniper Networks EX4300 Ethernet Switches:

- [Redundant Routing Engines on page 3](#)
- [Virtual Chassis on page 4](#)
- [VRRP on page 4](#)
- [Graceful Routing Engine Switchover on page 4](#)
- [Link Aggregation on page 5](#)
- [Nonstop Active Routing and Nonstop Bridging on page 5](#)

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 in case the master Routing Engine fails. 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 EX4300 Ethernet Switches configuring into a Virtual Chassis.

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, a simplified Layer 2 network topology that minimizes or eliminates the need for loop prevention protocols such as Spanning Tree Protocol (STP), and improved fault tolerance and high availability. A Virtual Chassis improves high availability for the following reasons:

- Dual Routing Engine support. A Virtual Chassis automatically has two Routing Engines—the switches in the master and backup **routing-engine** roles—and, therefore, provides more high availability options than standalone switches. Many high availability features are available for an EX Series Virtual Chassis that are not available on standalone EX Series switches.
- Increased fault tolerance. You increase your fault tolerance options when you configure your EX Series switches into a Virtual Chassis. You can, for instance, configure interfaces into a link aggregation group (LAG) with member interfaces on different member switches in the same Virtual Chassis to ensure network traffic is received by a Virtual Chassis even when a switch or physical interface in the Virtual Chassis fails.

You can configure up to ten EX4300 switches into an EX4300 Virtual Chassis. See *Understanding EX4300 Virtual Chassis*.

VRRP

You can configure Virtual Router Redundancy Protocol (VRRP) for IP and IPv6 on most switch interfaces, including Gigabit Ethernet interfaces, high-speed uplink interfaces, and logical interfaces. 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 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 GRES, 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 GRES 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 GRES 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, which increases high availability by ensuring that network traffic is received by the Virtual Chassis even if a single interface fails for any reason.

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 about LAGs, see *Understanding Aggregated Ethernet Interfaces and LACP*.

Nonstop Active Routing and Nonstop Bridging

Nonstop active routing (NSR) provides high availability in a switch or Virtual Chassis with redundant Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported Layer 3 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. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 3 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSR is enabled.

Nonstop bridging (NSB) provides the same mechanism for Layer 2 protocols. NSB 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. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 2 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSB is enabled.

To use NSR or NSB, you must also configure GRES.

**Related
Documentation**

- For more information about high availability features, see the [Junos OS High Availability Configuration Guide](#).
- *EX Series Virtual Chassis Overview*
- *Understanding EX4300 Virtual Chassis*
- [Understanding VRRP on EX Series Switches on page 27](#)
- *Understanding Aggregated Ethernet Interfaces and LACP*

CHAPTER 2

Nonstop Active Routing Overview

- [Understanding Nonstop Active Routing on EX Series Switches on page 7](#)

Understanding Nonstop Active Routing on EX Series Switches

You can configure nonstop active routing (NSR) on an EX Series switch with redundant Routing Engines or on an EX Series Virtual Chassis 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 107](#).



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 9](#).

**Related
Documentation**

- [Configuring Nonstop Active Routing on Switches on page 33](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 43](#)

CHAPTER 3

Nonstop Bridging Overview

- [Understanding Nonstop Bridging on EX Series Switches on page 9](#)

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 or on an EX Series Virtual Chassis with redundant Routing Engines.

NSB 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 NSB, see *EX Series Switch Software Features Overview* and *EX Series Virtual Chassis 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 7](#).

Related Documentation

- For information about configuring NSB on EX Series switches that do not support the Enhanced Layer 2 Software (ELS) CLI style, see [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)
- For information about configuring NSB on EX Series switches that support ELS, see [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 35](#)

Nonstop Software Upgrade Overview

- [Understanding Nonstop Software Upgrade on EX Series Switches on page 11](#)

Understanding Nonstop Software Upgrade on EX Series Switches

Nonstop software upgrade (NSSU) enables you to upgrade the software running on Juniper Networks EX Series Ethernet Switches with redundant Routing Engines and all member switches in EX Series Virtual Chassis by using a single command and with minimal network traffic disruption during the upgrade.

NSSU is supported on the following platforms:

- EX3300 Virtual Chassis
- EX4200 Virtual Chassis
- EX4300 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 network traffic—An NSSU minimizes network 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, EX4300, 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)*.

This topic covers:

- [Requirements for Performing an NSSU on page 12](#)
- [How an NSSU Works on page 14](#)
- [NSSU Limitations on page 17](#)
- [NSSU and Junos OS Release Support on page 17](#)
- [Overview of NSSU Configuration and Operation on page 18](#)

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)*.

- 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, EX4300, 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.



NOTE: For the EX4300 Virtual Chassis, you should enable the `vcp-no-hold-time` statement at the `[edit virtual-chassis]` hierarchy level before performing a software upgrade using NSSU. If you do not enable the `vcp-no-hold-time` statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see *Understanding Split and Merge in a Virtual Chassis*.

How an NSSU Works

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

- [EX3300, EX4200, EX4300, EX4500, and Mixed Virtual Chassis on page 14](#)
- [EX6200 and EX8200 Switches on page 14](#)
- [EX8200 Virtual Chassis on page 16](#)

EX3300, EX4200, EX4300, EX4500, and Mixed Virtual Chassis

When you request an NSSU on an EX3300, EX4200, EX4300, 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 (GRES) and nonstop active routing (NSR) 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 17](#) 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
EX4300 Virtual Chassis	13.2X51-D20 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 12](#). 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*.

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, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 36](#)
- *Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*
- [Configuring Nonstop Active Routing on Switches on page 33](#)
- *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*

CHAPTER 5

Power Management Overview

- [Understanding Power Management on EX Series Switches on page 19](#)

Understanding Power Management on EX Series Switches

The power management feature for Juniper Networks 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 20](#)
- [Power Supply Redundancy on page 23](#)

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 20](#)
- [Line Card Priority and Line Card Power on page 21](#)
- [Line Card Priority and PoE Power on page 21](#)
- [Line Card Priority and Changes in the Power Budget on page 22](#)

How a Line Card's Power Priority Is Determined

Using the CLI, you can assign an 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 on EX Series Switches (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 EX Series 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 24](#) 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)*
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 40](#)
- [Verifying Power Configuration and Use on page 77](#)

CHAPTER 6

VRRP Overview

- [Understanding VRRP on EX Series Switches on page 27](#)

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 on page 27](#)
- [Examples of VRRP Topologies on page 28](#)

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 2 on page 29.

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 1 on page 28 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 1: Basic VRRP on EX Series Switches

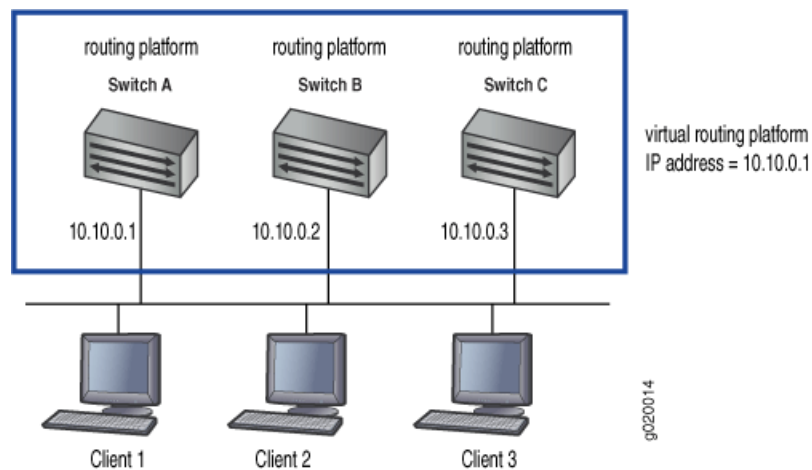
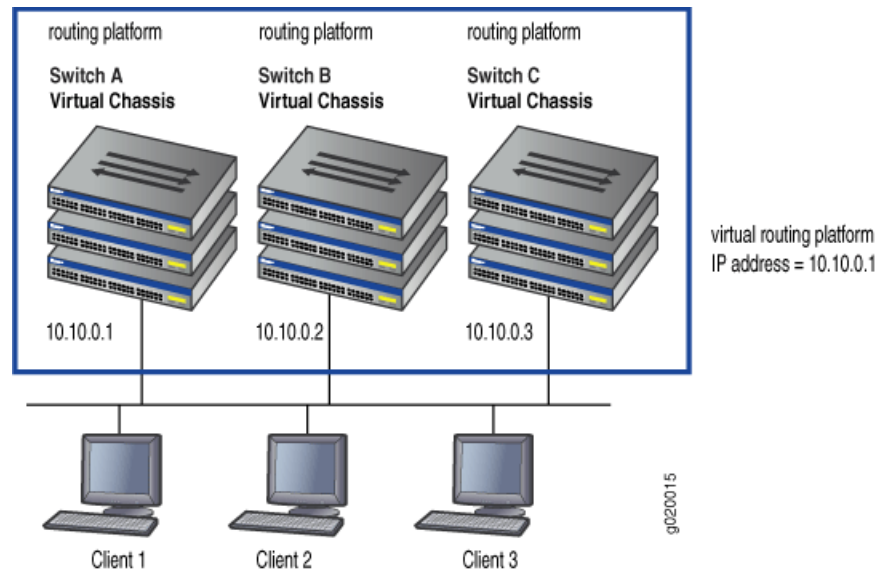


Figure 2 on page 29 illustrates a basic VRRP topology using Virtual Chassis configurations. Switch A, Switch B, and Switch C are each composed of multiple interconnected Juniper Networks EX4200 Ethernet Switches. Each Virtual Chassis configuration operates as a single switch, which is running VRRP, and 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 2: 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](#)
- [Configuring VRRP for IPv6 \(CLI Procedure\) on page 41](#)

PART 2

Configuration

- [Configuration Tasks on page 33](#)
- [Configuration Examples on page 43](#)
- [Configuration Statements on page 47](#)

CHAPTER 7

Configuration Tasks

- [Configuring Nonstop Active Routing on Switches on page 33](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 35](#)
- [Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 36](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 40](#)
- [Configuring VRRP for IPv6 \(CLI Procedure\) on page 41](#)

Configuring Nonstop Active Routing on Switches

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 NSR on an on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

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 issue the `commit synchronize` command, the candidate configuration is committed in the master Routing Engine. 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 Switches on page 43](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 107](#)
- [Understanding Nonstop Active Routing on EX Series Switches on page 7](#)
- [Nonstop Active Routing Concepts](#)

Configuring Nonstop Bridging on Switches (CLI Procedure)



NOTE: This task uses switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Limited support for NSB is also provided on QFX5100 and EX4600 standalone switches, but NSB is enabled *only* during an ISSU.

NSB 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. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 2 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSB is enabled.



NOTE: If you are using a QFX5100 or EX4600 standalone switch and you want to use ISSU, configure Graceful Routing Engine switchover (GRES), NSB and nonstop active routing (NSR). You must configure NSB, GRES, and NSR in order to run ISSU. However, GRES, NSB and NSR are enabled *only* during the upgrade. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.

To configure NSB:

1. Enable graceful Routing Engine switchover (GRES):


```
[edit chassis redundancy]
user@switch# set graceful-switchover
```
2. Enable NSB:


```
[edit protocols layer2-control]
user@switch# set nonstop-bridging
```
3. Synchronize configuration changes between the Routing Engines:


```
[edit system]
```

user@switch# **set commit** [synchronize](#)

If you try to commit a configuration that includes NSB 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 use the **commit synchronize** statement, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine 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.

**Related
Documentation**

- [Performing an In-Service Software Upgrade \(ISSU\)](#)
- [Understanding Nonstop Bridging on EX Series Switches on page 9](#)
- [Nonstop Bridging Concepts](#)
- [Understanding In-Service Software Upgrade \(ISSU\)](#)

[Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and 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 the following Virtual Chassis platforms:

- EX3300 Virtual Chassis
- EX4200 Virtual Chassis
- EX4300 Virtual Chassis
- EX4500 Virtual Chassis
- EX4550 Virtual Chassis
- All mixed Virtual Chassis composed of EX4200, EX4500, and EX4550 switches
- EX8200 Virtual Chassis

This topic covers:

- [Preparing the Switch for Software Installation on page 37](#)
- [Upgrading the Software Using NSSU on page 38](#)

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.

For information on configuring a preprovisioned Virtual Chassis, see *Configuring an EX3300 Virtual Chassis (CLI Procedure)*, *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*, and *Configuring an EX8200 Virtual Chassis (CLI Procedure)*.

- 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 Switches” on page 33](#) for information on how to enable it.

- For the EX4300 Virtual Chassis, you should enable the **vcp-no-hold-time** statement at the **[edit virtual-chassis]** hierarchy level before performing a software upgrade using

NSSU. If you do not enable the **vcp-no-hold-time** statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see *Understanding Split and Merge in a Virtual Chassis*.

- (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, EX4300 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 Nonstop Software Upgrade on EX Series Switches on page 11](#)
 - *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*
 - *Understanding Resilient Dual-Root Partitions on Switches*
 - *Understanding Software Installation on EX Series Switches*
 - *Troubleshooting Software Installation*
 - *Junos OS Package Names for EX Series Switches*
 - [Understanding Nonstop Software Upgrade on EX Series Switches on page 11](#)

Configuring Power Supply Redundancy (CLI Procedure)

By default, the power management feature in EX Series 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 \(CLI Procedure\)](#)
- [Verifying Power Configuration and Use on page 77](#)
- [Understanding Power Management on EX Series Switches on page 19](#)

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 96](#)
- [Understanding VRRP on EX Series Switches on page 27](#)

CHAPTER 8

Configuration Examples

- [Example: Configuring Nonstop Active Routing on Switches on page 43](#)

Example: Configuring Nonstop Active Routing on 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 or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

- [Requirements on page 43](#)
- [Overview and Topology on page 43](#)
- [Configuration on page 44](#)
- [Verification on page 45](#)
- [Troubleshooting on page 45](#)

Requirements

This example uses the following hardware and software components:

- An EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration
- Junos OS Release 10.4 or later for EX Series switches
- Junos OS Release 13.2X51-D20 or later for QFX Series switches

Overview and Topology

Configure nonstop active routing on any EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric 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 45](#)

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 45](#)

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 107](#).

**Related
Documentation**

- [Configuring Nonstop Active Routing on Switches on page 33](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 107](#)
- [Understanding Nonstop Active Routing on EX Series Switches on page 7](#)
- *Nonstop Active Routing Concepts*

CHAPTER 9

Configuration Statements

- [\[edit chassis\] Configuration Statement Hierarchy on EX Series Switches on page 48](#)
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 49](#)
- [\[edit protocols layer2-control\] Hierarchy Level on page 50](#)
- [\[edit protocols vrrp\] Hierarchy Level on page 51](#)
- [chassis on page 52](#)
- [synchronize on page 54](#)
- [failover \(Chassis\) on page 56](#)
- [fpc on page 57](#)
- [graceful-switchover on page 58](#)
- [hold-time on page 59](#)
- [inet6-advertise-interval on page 59](#)
- [keepalive-time on page 60](#)
- [n-plus-n \(Power Management\) on page 61](#)
- [nonstop-bridging on page 61](#)
- [nonstop-routing on page 62](#)
- [preempt on page 63](#)
- [priority on page 64](#)
- [psu on page 65](#)
- [redundancy \(Graceful Switchover\) on page 66](#)
- [redundancy \(Power Management\) on page 67](#)
- [traceoptions on page 68](#)
- [vcp-no-hold-time on page 71](#)
- [virtual-inet6-address on page 72](#)
- [virtual-link-local-address on page 73](#)
- [vrrp-inet6-group on page 74](#)

[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 [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit chassis\] Hierarchy Level on page 48](#)

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;
      lacp {
        link-protection non-revertive;
        system-priority system-priority-number
      }
    }
  }
  alarm {
    ethernet {
      link-down (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
  }
  container-devices {
    device-count device-count-number;
  }
  disk-partition {
    /config {
      level (full | high) {
        free-space (free-space-threshold-value | mb | percent);
      }
    }
  }
  /var {
    level (full | high) {
      free-space (free-space-threshold-value | mb | percent);
    }
  }
}
```

```

    }
  }
}
fpc slot-number {
  pic pic-number {
    no-multi-rate;
    q-pic-large-buffer (large-scale | small-scale);
  }
}
}
maximum-ecmp maximum-ecmp-routes;
lcd-menu {
  fpc slot-number {
    menu-item menu-name;
    disable;
  }
  pseudowire-service {
    device-count device-count-number;
  }
}
psu {
  redundancy {
    n-plus-n;
  }
  redundancy {
    graceful-switchover;
  }
}
slow-pfe-alarm;
}

```

Related Documentation

- [Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)
- [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 40](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)](#)

[\[edit interfaces\] 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](#)
- [\[edit interfaces et\] 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](#)
- [\[edit interfaces irb\] 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 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 Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- *Configuring Integrated Routing and Bridging 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 layer2-control] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  layer2-control {
    bpdu-block {
      disable-timeout seconds;
      interface [ interface-names ];
    }
    mac-rewrite {
      interface interface-name {
        enable-all-ifl;
        protocol {
          cdp;
          stp;
          vtp;
          pvstp;
        }
      }
    }
  }
  nonstop-bridging;
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
}
```

Related Documentation

- *layer2-control*

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols vrrp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  vrrp {
    asymmetric-hold-time;
    delegate-processing;
    failover-delay milliseconds;
    global-advertisements-threshold advertisement-value;
    skew-timer-disable;
    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;
    }
    version-3;
  }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*
- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

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 by 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 40](#)
 - *Configuring the Power Priority of Line Cards (CLI Procedure)*
 - *Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure)*

synchronize

Syntax	synchronize;
Hierarchy Level	[edit system commit]
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 perform 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: If you configure the **commit synchronize** statement at the [edit system] hierarchy level and issue a **commit** in the master Routing Engine, the master configuration is automatically synchronized with the backup. However, if the backup Routing Engine is down when you issue the **commit**, the Junos OS displays a warning and commits the candidate configuration in the master Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the master. A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.



NOTE: When you configure nonstop active routing (NSR), you must configure 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. 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.

On the TX Matrix Plus router, synchronization only occurs between the Routing Engines within the switch-fabric chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the line-card chassis (LCC). That is, the master Routing Engine on the TX Matrix Plus router distributes the configuration to the master Routing Engine on each LCC. Likewise, the backup Routing Engine on the TX Matrix Plus router distributes the configuration to the backup Routing Engine on each LCC.

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) Quit configuration mode if the commit synchronization succeeds.

at—(Optional) Time at which to activate configuration changes.

comment—(Optional) Write a message to the commit log.

force—(Optional) Force a commit synchronization on the other Routing Engine (ignore warnings).

scripts—(Optional) Push scripts to the other Routing Engine.

Required Privilege **system**—To view this statement in the configuration.

Level **system-control**—To add this statement to the configuration.

Related • *Synchronizing the Routing Engine Configuration*
Documentation • *Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically*

failover (Chassis)

Syntax	<pre>failover { on-disk-failure; on-loss-of-keepalives; }</pre>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Specify conditions on the master Routing Engine that cause the backup router to take mastership.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• graceful-switchover on page 58• <i>On Detection of a Hard Disk Error on the Master Routing Engine</i>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>• <i>High Availability Features for EX Series Switches Overview</i>

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.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description Specify the port of the SFP+ uplink module for which you want to configure the operating mode, or 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 power priority for the slot, your change does not take 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 along 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, EX4500, and EX4550 switches. The FPC value refers to the switch itself.
- 0–9—EX4200, EX4500, or EX4550 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	<ul style="list-style-type: none">• <i>Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure)</i>• <i>Configuring the Power Priority of Line Cards (CLI Procedure)</i>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i>

graceful-switchover

Syntax	graceful-switchover;
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	For switches with more than one Routing Engine, including those in a Virtual Chassis or a Virtual Chassis Fabric, 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	<ul style="list-style-type: none">• Example: Configuring Nonstop Active Routing on Switches on page 43• <i>Configuring Graceful Routing Engine Switchover</i>• <i>Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)</i>• Configuring Nonstop Active Routing on Switches on page 33• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>

hold-time

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> preempt]
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	<i>seconds</i> —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	<ul style="list-style-type: none"> • Configuring VRRP for IPv6 (CLI Procedure) on page 41

inet6-advertise-interval

Syntax	<code>inet6-advertise-interval milliseconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]
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	<i>milliseconds</i> —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	<ul style="list-style-type: none"> • Configuring VRRP for IPv6 (CLI Procedure) on page 41

keepalive-time

Syntax	<code>keepalive-time <i>seconds</i>;</code>
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	<p>The on-loss-of-keepalives statement at the [edit chassis redundancy failover] hierarchy level must be included for failover to occur.</p> <p>When the on-loss-of-keepalives statement is included and graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the on-loss-of-keepalives statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds.</p>
Options	<i>seconds</i> —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	<ul style="list-style-type: none">• failover on page 56• graceful-switchover on page 58• <i>on-loss-of-keepalives</i>• <i>High Availability Features for EX Series Switches Overview</i>

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	<ul style="list-style-type: none"> • Configuring Power Supply Redundancy (CLI Procedure) on page 40

nonstop-bridging

Syntax	nonstop-bridging;
Hierarchy Level	[edit protocols layer2-control]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	For platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Synchronizing the Routing Engine Configuration</i> • <i>Configuring Nonstop Bridging</i> • For information about configuring NSB on EX Series switches that do not support the Enhanced Layer 2 Software (ELS) CLI style, see <i>Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)</i> • For information about configuring NSB on switches that support ELS, see Configuring Nonstop Bridging on Switches (CLI Procedure) on page 35

nonstop-routing

Syntax nonstop-routing;

Hierarchy Level [edit routing-options]




NOTE: Although nonstop-routing is also a valid keyword at the logical-systems hierarchy level, it is not supported.

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. Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches
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	<ul style="list-style-type: none">• <i>Configuring Nonstop Active Routing</i>

preempt

Syntax	(preempt no-preempt) { hold-time seconds; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	<p>Configure whether a backup router can preempt a master router:</p> <ul style="list-style-type: none"> • preempt—Allow the master router to be preempted. • no-preempt—Prohibit the preemption of the master router. <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring VRRP for IPv6 (CLI Procedure) on page 41

priority

Syntax	<code>priority number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
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)
<div> NOTE: Priority 255 cannot be assigned to routed VLAN interfaces (RVIs).</div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP for IPv6 (CLI Procedure) on page 41

psu

Syntax	<pre>psu { redundancy { n-plus-n (Power Management); } }</pre>
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	<p>Configure <i>N+N</i> power supply redundancy for power management on an EX6200 or EX8200 switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Power Supply Redundancy (CLI Procedure) on page 40

redundancy (Graceful Switchover)

Syntax	<pre>redundancy { failover { on-disk-failure; on-loss-of-keepalives; } graceful-switchover; }</pre>
Hierarchy Level	[edit chassis]
Release Information	<p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately.</p>
Default	Redundancy is enabled for the Routing Engines.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• graceful-switchover on page 58• <i>Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)</i>• <i>Configuring Graceful Routing Engine Switchover</i>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>• <i>High Availability Features for EX Series Switches Overview</i>

redundancy (Power Management)

Syntax	<code>redundancy { n-plus-n (Power Management); }</code>
Hierarchy Level	[edit <code>chassis psu</code>]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	<p>Configure $N+N$ power supply redundancy for power management on an EX6200 or EX8200 switch.</p> <p>The remaining statement is explained separately.</p>
Default	$N+1$ power supply redundancy is configured by default.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Power Supply Redundancy (CLI Procedure) on page 40

traceoptions

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>Values:</p> <p>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.</p> <p>file <i>filename</i>—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.</p>

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**—Nonstop active routing synchronization
- **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: **xk** to specify KB, **xm** to specify MB, or **xg** 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	routing and trace—To view this statement in the configuration.
Level	routing-control and trace-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Example: Tracing Global Routing Protocol Operations</i>
------------------------------	--

vcp-no-hold-time

Syntax	vcp-no-hold-time;
Hierarchy Level	[edit virtual-chassis]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Description	<p>Disable the Virtual Chassis port (VCP) holddown timer for all VCPs in the Virtual Chassis or Virtual Chassis Fabric (VCF).</p> <p>The VCP holddown timer is an internal mechanism that delays a Virtual Chassis reconvergence for several seconds when a VCP becomes inactive. The purpose of this delay is to provide the VCP time to return online without having to reconverge the Virtual Chassis to adjust to the inactive VCP. All traffic to the VCP is dropped while the VCP is inactive. If the VCP remains down for a time that exceeds the VCP holddown timer, a Virtual Chassis reconvergence occurs.</p> <p>When this statement is enabled, the VCP holddown timer is disabled and the Virtual Chassis reconvergence occurs when a VCP becomes inactive. The period of time where traffic is dropped waiting for the VCP to return online is avoided.</p> <p>We recommend enabling this statement after a Virtual Chassis is operational. We recommend disabling this statement when you are adding or removing member switches from your Virtual Chassis.</p> <p>The VCP holddown timer cannot be viewed and is not user-configurable. You can only control whether the VCP holddown timer is enabled or disabled by configuring this statement.</p>
	<p> NOTE: For the EX4300 Virtual Chassis, you should enable the <code>vcp-no-hold-time</code> statement before performing a software upgrade using NSSU. If you do not enable the <code>vcp-no-hold-time</code> statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see <i>Understanding Split and Merge in a Virtual Chassis</i>.</p>
Default	The VCP holddown timer is enabled by default on all devices that support this statement.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding EX4300 Virtual Chassis</i> • <i>Understanding QFX Series Virtual Chassis</i>

- [Understanding EX Series Virtual Chassis Components](#)
- [Understanding QFX Series Virtual Chassis Components](#)

virtual-inet6-address

Syntax	virtual-inet6-address [<i>addresses</i>];
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
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	<i>addresses</i> —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	<ul style="list-style-type: none">• Configuring VRRP for IPv6 (CLI Procedure) on page 41

virtual-link-local-address

Syntax	<code>virtual-link-local-address <i>ipv6-address</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-inet6-group <i>group-id</i>]</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]</p>
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	<i>ipv6-address</i> —Virtual link local IPv6 address for VRRP for an IPv6 group.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring VRRP for IPv6 (CLI Procedure) on page 41

vrrp-inet6-group

Syntax	<pre>vrrp-inet6-group <i>group-id</i> { inet6-advertise-interval <i>milliseconds</i>; preempt { hold-time <i>seconds</i>; } priority <i>number</i>; virtual-inet6-address; virtual-link-local-address }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i>]
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	<p>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. The MAC address 00-00-5E-00-02-{VRID} is reserved for VRRP, as defined in RFC 5798. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>Range: 0 through 255</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP for IPv6 (CLI Procedure) on page 41

PART 3

Administration

- [Verification Tasks on page 77](#)
- [Operational Commands on page 81](#)

CHAPTER 10

Verification Tasks

- [Verifying Power Configuration and Use on page 77](#)

Verifying Power Configuration and Use

Purpose Verify on an EX Series 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
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 40](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)

CHAPTER 11

Operational Commands

- request system software nonstop-upgrade
- show chassis power-budget-statistics
- show vrrp

request system software nonstop-upgrade

Syntax	<code>request system software nonstop-upgrade (<i>package-name</i> set [<i>package-name</i> <i>package-name</i>])</code> <code><force-host></code> <code><no-copy></code> <code><no-old-master-upgrade></code> <code><reboot></code> <code><unlink></code>
Release Information	Command introduced in Junos OS Release 10.4 for EX Series switches. Option set [<i>package-name package-name</i>] added in Junos OS Release 12.1 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D20 for the QFX Series.
Description	<p>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:</p> <ul style="list-style-type: none">• When you execute this command on an EX3300, EX4200, EX4300, EX4500, or EX4550 Virtual Chassis or QFX3500 and QFX3600 Virtual Chassis, a fixed configuration of switches in a Virtual Chassis Fabric (QFX3500/QFX3600 and QFX5100 switches) or for a mixed Virtual Chassis Fabric composed of any combination of QFX3500/QFX3600, QFX5100, and EX4300 switches, or a mixed Virtual Chassis composed of any combination of EX4200, EX4500, and EX4550 switches, all members are upgraded. The original Virtual Chassis backup becomes the master. The original master is automatically upgraded and rebooted and rejoins the Virtual Chassis as the backup after the upgrade completes.• When you execute this command on an EX6200 or EX8200 switch, both the backup and master Routing Engines are upgraded, with the original backup Routing Engine becoming the new master at the end of the upgrade. 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. <p>This command has the following requirements:</p> <ul style="list-style-type: none">• 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)*.

- 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, EX4300, EX4500, EX4550, QFX3500 and QFX3600 Virtual Chassis, and mixed Virtual Chassis, and Virtual Chassis Fabric) or on different line cards (for EX6200 and EX8200 switches, and for EX8200 Virtual Chassis).
- For EX3300, EX4200, EX4300, EX4500, EX4550, QFX3500 and QFX3600 Virtual Chassis, and mixed Virtual Chassis:
 - The Virtual Chassis members must be connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
 - The Virtual Chassis master and backup must be adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in 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.
- For Virtual Chassis Fabric:
 - Only two pre-provisioned members in the routing engine role are supported. If more than two routing engines are configured, a warning will be issued, and NSSU will stop.
 - The Virtual Chassis Fabric members are connected in a spine and leaf topology. A spine and leaf topology prevents the Virtual Chassis Fabric from splitting during an NSSU. Each leaf device must be connected to both spine devices.
 - The Virtual Chassis Fabric must be preprovisioned so that the line card role has been explicitly assigned to member switches acting in a line card role, and that the routing engine role has been explicitly assigned to member switches acting in a routing engine role. During an NSSU, the Virtual Chassis Fabric members must maintain their roles—the master and backup must maintain their master and backup roles (although

mastership will change), the member switches must remain their routing engine roles, and the remaining switches must maintain their linecard roles.

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

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

- */var/tmp/package-name*—For a software package or bundle that is being installed from a local directory on the switch.
- *protocol://hostname/pathname/package-name*—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *protocol* with one of the following:
 - **ftp**—File Transfer Protocol.
Use *ftp://hostname/pathname/package-name*. To specify authentication credentials, use *ftp://<username>:<password>@hostname/pathname/package-name*. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
 - **http**—Hypertext Transfer Protocol.
Use *http://hostname/pathname/package-name*. To specify authentication credentials, use *http://<username>:<password>@hostname/pathname/package-name*. If a password is required and you omit it, you are prompted for it.
 - **scp**—Secure copy (available only for Canada and U.S. version).
Use *scp://hostname/pathname/package-name*. To specify authentication credentials, use *scp://<username>:<password>@hostname/pathname/package-name*.



NOTE: The *pathname* in the protocol is the relative path to the user home directory on the remote system and not the root directory.

set [*package-name package-name*]—(Mixed Virtual Chassis only) Locations of the EX4200 and the EX4500 installation packages. These packages must be for the same Junos OS release. See the description of the *package-name* option for information about how to specify the location of the installation packages.

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

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

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

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



NOTE: If you do not use the **reboot** option on an EX8200 Virtual Chassis, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module to perform the manual reboot of the backup Routing Engines.

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

Required Privilege Level maintenance

Related Documentation

- *show chassis nonstop-upgrade*
- [Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 36](#)
- *Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade*
- *Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*

List of Sample Output

[request system software nonstop-upgrade \(EX4200 Virtual Chassis\) on page 86](#)
[request system software nonstop-upgrade \(EX6200 Switch\) on page 87](#)
[request system software nonstop-upgrade reboot \(EX8200 Switch\) on page 88](#)
[request system software nonstop-upgrade no-old-master-upgrade \(EX8200 Switch\) on page 89](#)
[request system software nonstop-upgrade reboot \(EX8200 Virtual Chassis\) on page 89](#)

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
  FPC 3     Online (ISSU)
  Offlined by CLI command
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 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 EX Series switch.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Verifying Power Configuration and Use on page 77 • Configuring the Power Priority of Line Cards (CLI Procedure) • Configuring Power Supply Redundancy (CLI Procedure) on page 40
List of Sample Output	show chassis power-budget-statistics (EX6200 Switch) on page 94 show chassis power-budget-statistics (EX8200 Switch) on page 94
Output Fields	Table 5 on page 92 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 5: show chassis power-budget-statistics Output Fields

Field Name	Field Description
<i>PSU n (supply type)</i>	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 <i>N+1</i> or <i>N+N</i> .
Base power reserved	Total number of watts reserved for the switch.
Non-PoE power being consumed	The amount of power, in W, currently being consumed for PoE.
Power Reserved for the Chassis	<p>Power reserved for the chassis:</p> <ul style="list-style-type: none"> • For an EX6200 switch, 500 W. • For an EX8208 switch: 1600 W in an <i>N+1</i> configuration; 1200 W in an <i>N+N</i> configuration • For an EX8216 switch: 2400 W in an <i>N+1</i> configuration; 1800 W in an <i>N+N</i> configuration <p>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.</p>

Table 5: show chassis power-budget-statistics Output Fields (*continued*)

Field Name	Field Description
Fan Tray Statistics	<p>(EX6200 switch only) Information about the fan tray:</p> <ul style="list-style-type: none"> • Base power—Power allocated to the fan tray in the power budget. This allocation is included in Power Reserved for the Chassis. • 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 <i>n</i> (<i>card type</i>)	<p>Information about the line card installed in slot <i>n</i>. For EX6200 switches, information about the SRE modules in slot 4 and slot 5 is also shown.</p> <ul style="list-style-type: none"> • Base power—For line cards without PoE ports, the total power allocated to the line card. 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. • 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. • 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. For line cards without PoE ports, the value is always 0 W. • 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.
Total PoE power consumed	The amount of power that has been consumed by PoE.
Total PoE power remaining	The amount of available power remaining that can be used for 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 <i>N</i> +1 or <i>N</i> + <i>N</i> redundancy requirements, this value is 0. PoE power allocations are not included in the calculation of this value.

Table 5: 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
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 vrrp

Syntax	<pre>show vrrp <brief detail extensive summary> <interface <i>interface-name</i>> <track interfaces></pre>
Release Information	<p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	Display information and status about VRRP groups.
Options	<p>none—(Same as brief) Display brief status information about all VRRP interfaces.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display information and status about the specified VRRP interface.</p> <p>track interfaces—(Optional) Display information and status about VRRP track interfaces.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Configuring VRRP for IPv6 (CLI Procedure) on page 41
List of Sample Output	<p>show vrrp on page 101</p> <p>show vrrp brief on page 101</p> <p>show vrrp detail (IPv6) on page 101</p> <p>show vrrp detail (Route Track) on page 102</p> <p>show vrrp extensive on page 102</p> <p>show vrrp interface on page 104</p> <p>show vrrp summary on page 104</p> <p>show vrrp track detail on page 104</p> <p>show vrrp track summary on page 105</p>
Output Fields	<p>Table 6 on page 96 lists the output fields for the show vrrp command. Output fields are listed in the approximate order in which they appear.</p>

Table 6: 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 6: 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: <ul style="list-style-type: none"> • Advertisement sent—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. • Advertisement received—Number of VRRP advertisement PDUs received by the interface. • Packets received—Number of VRRP packets received for VRRP groups on the interface. • 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: <ul style="list-style-type: none"> • Invalid IPAH next type received—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. • Invalid VRRP ttl value received—Number of packets received whose IP time-to-live (TTL) value is not 255. • Invalid VRRP version received—Number of packets received whose VRRP version is not 2. • Invalid VRRP pdu type received—Number of packets received whose VRRP PDU type is not 1. • Invalid VRRP authentication type received—Number of packets received whose VRRP authentication is not none, simple, or md5. • Invalid VRRP IP count received—Number of packets received whose VRRP IP count exceeds 8. • 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 6: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type and Address	Identifier for the address and the address itself: <ul style="list-style-type: none"> • lcl—Configured local interface address. • mas—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router. • vip—Configured virtual IP addresses. 	none, brief, summary
Interface state or Int state	State of the physical interface: <ul style="list-style-type: none"> • down—The device is present and the link is unavailable. • not present—The interface is configured, but no physical device is present. • unknown—The VRRP process has not had time to query the kernel about the state of the interface. • 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: <ul style="list-style-type: none"> • backup—The interface is acting as the backup router interface. • bringup—VRRP is just starting, and the physical device is not yet present. • 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. • initializing—VRRP is initializing. • master—The interface is acting as the master router interface. • 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 6: 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. 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 6: 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: <ul style="list-style-type: none"> • 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. • 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. • 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. • 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. • 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. • 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: <ul style="list-style-type: none"> • Idle to master transitions—Number of times that the VRRP instance transitioned from the idle state to the master state. • Idle to backup transitions—Number of times that the VRRP instance transitioned from the idle state to the backup state. • Backup to master transitions—Number of times that the VRRP instance transitioned from the backup state to the master state. • 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 6: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
VR state	VRRP information: <ul style="list-style-type: none"> • backup—The interface is acting as the backup router interface. • bringup—VRRP is just starting, and the physical device is not yet present. • 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. • initializing—VRRP is initializing. • master—The interface is acting as the master router interface. • transition—The interface is changing between being the backup and being the master router. 	none, brief
Timer	VRRP timer information: <ul style="list-style-type: none"> • A—Time, in seconds, until the advertisement timer expires. • 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 1c1
ge2001:db8::12:1:1:1
ge2001:db8:0:1:12:1:1:99
                                     vip
                                     vip
ge2001:db8::12:1:1:99
ge-0/0/2.131   up         1      master    A 0.364 1c1
ge2001:db8::13:1:1:1
                                     vip
ge2001:db8:0:1:13:1:1:99
                                     vip
ge2001:db8::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 101](#).

show vrrp detail (IPv6)

```

user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address:
ge2001:db8::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: ge2001:db8:0:1:12:1:1:99,
ge2001:db8::12:1:1:99
  Advertisement timer: 1.121s, Master router: ge2001:db8:0:1:12:1:1:1
  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:
ge2001:db8::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: ge2001:db8:0:1:13:1:1:99,
ge2001:db8::13:1:1:99
Advertisement timer: 0.327s, Master router: ge2001:db8:0:1:13:1:1:1
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: 192.0.2.30/24
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: 192.0.2.100
Advertisement timer: 1.218s, Master router: 192.0.2.30
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
    198.51.100.0/24  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:
ge2001:db8::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: ge2001:db8:0:1:12:1:1:99,  
ge2001:db8::12:1:1:99
Advertisement timer: 1.034s, Master router: ge2001:db8:0:1:12:1:1:1
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:
ge2001:db8::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: ge2001:db8:0:1:13:1:1:99,
ge2001:db8::13:1:1:99
  Advertisement timer: 0.396s, Master router: ge2001:db8:0:1:13:1:1:1
  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:
ge2001:db8::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: ge2001:db8:0:1:12:1:1:99,
gec2001:db8::12:1:1:99
  Advertisement timer: 0.789s, Master router: ge2001:db8:0:1:12:1:1:1
  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
```


Troubleshooting Procedures

- [Tracing Nonstop Active Routing Synchronization Events on page 107](#)

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 Switches on page 33](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 43](#)
- [Example: Configuring Nonstop Active Routing](#)