



Junos[®] OS for EX Series Ethernet Switches

High Availability Feature Guide for EX4600 Switches

Release
15.1



Modified: 2016-07-19

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.

Junos[®] OS for EX Series Ethernet Switches High Availability Feature Guide for EX4600 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	Graceful Restart	
Chapter 1	Using Graceful Restart	3
	Graceful Restart Concepts	3
	Configuring Routing Protocols Graceful Restart	4
	Enabling Graceful Restart	4
	Configuring Graceful Restart Options for BGP	5
	Configuring Graceful Restart Options for ES-IS	6
	Configuring Graceful Restart Options for IS-IS	6
	Configuring Graceful Restart Options for OSPF and OSPFv3	7
	Configuring Graceful Restart Options for RIP and RIPng	9
	Configuring Graceful Restart Options for PIM Sparse Mode	9
	Tracking Graceful Restart Events	10
Part 2	Graceful Routing Engine Switchover	
Chapter 2	Using GRES	13
	Understanding Graceful Routing Engine Switchover	13
	Graceful Routing Engine Switchover Concepts	13
	Effects of a Routing Engine Switchover	18
	Graceful Routing Engine Switchover System Requirements	19
	Graceful Routing Engine Switchover Platform Support	19
	Graceful Routing Engine Switchover Feature Support	20
	Graceful Routing Engine Switchover DPC Support	22
	Graceful Routing Engine Switchover and Subscriber Access	22
	Graceful Routing Engine Switchover PIC Support	22
	Configuring Graceful Routing Engine Switchover	23
	Enabling Graceful Routing Engine Switchover	23
	Configuring Graceful Routing Engine Switchover with Graceful Restart	23

	Synchronizing the Routing Engine Configuration	23
	Verifying Graceful Routing Engine Switchover Operation	25
	Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)	25
	Resetting Local Statistics	26
Part 3	Nonstop Bridging and Routing	
Chapter 3	Using Nonstop Bridging and Routing	29
	Nonstop Active Routing Concepts	29
	Nonstop Active Routing System Requirements	32
	Nonstop Active Routing Platform and Switching Platform Support	32
	Nonstop Active Routing Protocol and Feature Support	33
	Nonstop Active Routing BFD Support	36
	Nonstop Active Routing BGP Support	37
	Nonstop Active Routing Layer 2 Circuit and VPLS Support	38
	Nonstop Active Routing PIM Support	38
	Nonstop Active Routing MSDP Support	41
	Nonstop Active Routing Support for RSVP-TE LSPs	41
	Nonstop Bridging Concepts	43
	Nonstop Bridging System Requirements	45
	Platform Support	45
	Protocol Support	46
	Configuring Nonstop Active Routing on Switches	46
	Tracing Nonstop Active Routing Synchronization Events	47
	Configuring Nonstop Bridging on Switches (CLI Procedure)	49
	Resetting Local Statistics	50
	Example: Configuring Nonstop Active Routing on Switches	51
Part 4	Virtual Router Redundancy Protocol	
Chapter 4	Using VRRP	57
	Understanding VRRP	57
	Overview of VRRP	57
	Sample VRRP Topology	58
	Configuring Basic VRRP Support for QFX	60
	Configuring VRRP Authentication (IPv4 Only)	61
	Configuring the Startup Period for VRRP Operations	62
	Configuring the Advertisement Interval for the VRRP Master	62
	Modifying the Advertisement Interval in Seconds	62
	Modifying the Advertisement Interval in Milliseconds	63
	Configuring VRRP Preemption and Hold Time	63
	Configuring VRRP Preemption	63
	Configuring the Preemption Hold Time	64
	Overriding the Hold Time	64
	Configuring a Route to Be Tracked	64
	Configuring a Logical Interface to Be Tracked	65
	Configuring a Backup to Accept Packets Destined for the Virtual IP Address	67
	Configuring Passive ARP Learning for VRRP Backups	67

	Configuring the Silent Period	68
	Configuring Inheritance for a VRRP Group	68
	Example: Configuring VRRP for Load Sharing	69
	Troubleshooting VRRP	74
Part 5	Nonstop Software Upgrade	
Chapter 5	Using NSSU	79
	Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade	79
	Preparing the Switch for Software Installation	79
	Upgrading the Software Using NSSU	80
Part 6	Configuration Statements and Operational Commands	
Chapter 6	Configuration Statements for Graceful Restart	85
	disable	86
	disable (BGP Graceful Restart)	87
	graceful-restart (Enabling Globally)	88
	graceful-restart (Protocols BGP)	90
	graceful-restart	92
	helper-disable (OSPF)	94
	no-strict-lsa-checking	95
	notify-duration	96
	redundancy (Graceful Switchover)	97
	restart-duration	98
	restart-time (BGP Graceful Restart)	99
	stale-routes-time	100
Chapter 7	Configuration Statements for Graceful Switchover	101
	graceful-switchover	101
	redundancy (Graceful Switchover)	102
Chapter 8	Configuration Statements for Nonstop Bridging and Routing	103
	nonstop-bridging	103
	nonstop-routing	104
	synchronize	105
	traceoptions	107
Chapter 9	Configuration Statements for VRRP	111
	accept-data	112
	advertise-interval	113
	asymmetric-hold-time	114
	authentication-key	115
	authentication-type	116
	bandwidth-threshold	117
	failover-delay	118
	fast-interval	119
	hold-time (VRRP)	120
	interface	121
	preempt (VRRP)	122

	priority (Protocols VRRP)	123
	priority-cost (VRRP)	124
	priority-hold-time	125
	route (Interfaces)	126
	startup-silent-period	127
	traceoptions	128
	track (VRRP)	130
	virtual-address	131
	vrrp-group	132
Chapter 10	Operational Commands for Graceful Restart	135
	Verifying Graceful Restart Operation	135
	Graceful Restart Operational Mode Commands	135
	Verifying BGP Graceful Restart	136
	Verifying IS-IS and OSPF Graceful Restart	136
	Verifying CCC and TCC Graceful Restart	137
	show bgp neighbor	138
	show log	155
	show (ospf ospf3) overview	159
Chapter 11	Operational Commands for Graceful Switchover	165
	show system switchover	166
	show task replication	172
Chapter 12	Operational Command for Nonstop Routing	175
	show task replication	176
Chapter 13	Operational Command for NSSU	179
	request system software nonstop-upgrade	180
Chapter 14	Operational Commands for VRRP	191
	show vrrp	192

List of Figures

Part 2	Graceful Routing Engine Switchover	
Chapter 2	Using GRES	13
	Figure 1: Preparing for a Graceful Routing Engine Switchover	16
	Figure 2: Graceful Routing Engine Switchover Process	17
Part 3	Nonstop Bridging and Routing	
Chapter 3	Using Nonstop Bridging and Routing	29
	Figure 3: Nonstop Active Routing Switchover Preparation Process	30
	Figure 4: Nonstop Active Routing During a Switchover	31
	Figure 5: Nonstop Bridging Switchover Preparation Process	44
	Figure 6: Nonstop Bridging During a Switchover	45
Part 4	Virtual Router Redundancy Protocol	
Chapter 4	Using VRRP	57
	Figure 7: Basic VRRP Topology	59
	Figure 8: VRRP Load-Sharing Configuration	70

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiii
Part 2	Graceful Routing Engine Switchover	
Chapter 2	Using GRES	13
	Table 3: Effects of a Routing Engine Switchover	18
	Table 4: Graceful Routing Engine Switchover Feature Support	20
Part 3	Nonstop Bridging and Routing	
Chapter 3	Using Nonstop Bridging and Routing	29
	Table 5: Nonstop Active Routing Platform Support	32
	Table 6: Nonstop Active Routing Protocol and Feature Support	34
Part 4	Virtual Router Redundancy Protocol	
Chapter 4	Using VRRP	57
	Table 7: Interface State and Priority Cost Usage	66
	Table 8: Settings for VRRP Load-Sharing Example	71
Part 6	Configuration Statements and Operational Commands	
Chapter 10	Operational Commands for Graceful Restart	135
	Table 9: show bgp neighbor Output Fields	139
	Table 10: show ospf overview Output Fields	160
Chapter 11	Operational Commands for Graceful Switchover	165
	Table 11: show system switchover Output Fields	168
	Table 12: show task replication Output Fields	172
Chapter 12	Operational Command for Nonstop Routing	175
	Table 13: show task replication Output Fields	176
Chapter 14	Operational Commands for VRRP	191
	Table 14: show vrrp Output Fields	192

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.	<code>user@host> show chassis alarms</code> <code>No alarms currently active</code>
<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

Graceful Restart

- [Using Graceful Restart on page 3](#)

CHAPTER 1

Using Graceful Restart

- [Graceful Restart Concepts on page 3](#)
- [Configuring Routing Protocols Graceful Restart on page 4](#)

Graceful Restart Concepts

With routing protocols, any service interruption requires that an affected router recalculate adjacencies with neighboring routers, restore routing table entries, and update other protocol-specific information. An unprotected restart of a router can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Three main types of graceful restart are available on Juniper Networks routing platforms:

- Graceful restart for aggregate and static routes and for routing protocols—Provides protection for aggregate and static routes and for Border Gateway Protocol (BGP), End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), next-generation RIP (RIPng), and Protocol Independent Multicast (PIM) sparse mode routing protocols.
- Graceful restart for MPLS-related protocols—Provides protection for Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), circuit cross-connect (CCC), and translational cross-connect (TCC). (Not supported on OCX Series switches.)
- Graceful restart for virtual private networks (VPNs)—Provides protection for Layer 2 and Layer 3 VPNs.

Graceful restart works similarly for routing protocols and MPLS protocols and combines components of these protocol types to enable graceful restart in VPNs. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart thus enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Most graceful restart implementations define two types of routers—the restarting router and the helper router. The restarting router requires rapid restoration of forwarding state

information so it can resume the forwarding of network traffic. The helper router assists the restarting router in this process. Graceful restart configuration statements typically affect either the restarting router or the helper router.

**Related
Documentation**

- *[Understanding High Availability Features on Juniper Networks Routers](#)*
- *[Graceful Restart System Requirements](#)*
- *[Graceful Restart for Aggregate and Static Routes](#)*
- *[Graceful Restart and Routing Protocols](#)*
- *[Graceful Restart and MPLS-Related Protocols](#)*
- *[Graceful Restart and Layer 2 and Layer 3 VPNs](#)*
- *[Graceful Restart on Logical Systems](#)*
- *[Configuring Graceful Restart](#)*
- *[Configuring Graceful Restart for QFabric Systems](#)*

Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- [Enabling Graceful Restart on page 4](#)
- [Configuring Graceful Restart Options for BGP on page 5](#)
- [Configuring Graceful Restart Options for ES-IS on page 6](#)
- [Configuring Graceful Restart Options for IS-IS on page 6](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 7](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 9](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 9](#)
- [Tracking Graceful Restart Events on page 10](#)

Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {  
    graceful-restart;  
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



NOTE: Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



NOTE: To set BGP graceful restart properties or disable them for a group, include the desired statements at the `[edit protocols bgp group group-name graceful-restart]` hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the `[edit protocols bgp group group-name neighbor ip-address graceful-restart]` hierarchy level.



NOTE: Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

Configuring Graceful Restart Options for ES-IS

On J Series Services Routers, to configure the duration of the ES-IS graceful restart period, include the **restart-duration** statement at the `[edit protocols esis graceful-restart]` hierarchy level.

```
[edit]
protocols {
  esis {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable ES-IS graceful restart capability, include the **disable** statement at the `[edit protocols esis graceful-restart]` hierarchy level.

Configuring Graceful Restart Options for IS-IS

To configure the duration of the IS-IS graceful restart period, include the **restart-duration** statement at the `[edit protocols isis graceful-restart]` hierarchy level.

```
[edit]
protocols {
  isis {
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

```
}
```

To disable IS-IS graceful restart helper capability, include the **helper-disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level. To disable IS-IS graceful restart capability, include the **disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level.



NOTE: If you configure Bidirectional Forwarding Detection (BFD) and graceful restart for IS-IS, graceful restart might not work as expected.



NOTE: If adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the hold-time for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.



NOTE: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols isis]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 10](#).

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  ospf | ospfv3 {
    graceful-restart {
      disable;
      helper-disable
      no-strict-lsa-checking;
      notify-duration seconds;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
  graceful-restart {
    helper-disable <both | restart-signaling | standard>
  }
```

To reenble the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.



NOTE:

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the **[edit protocols ospfv3 graceful-restart]** hierarchy level.



TIP: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols (ospf | ospf3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 10](#).



NOTE: You cannot enable OSPFv3 graceful restart between a routing platform running Junos OS Release 7.5 and earlier and a routing platform running Junos OS Release 7.6 or later. As a workaround, make sure both routing platforms use the same Junos OS version.



NOTE: If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

Configuring Graceful Restart Options for RIP and RIPng

To configure the duration of the RIP or RIPng graceful restart period, include the **restart-time** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  (rip | ripng) {
    graceful-restart {
      disable;
      restart-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable RIP or RIPng graceful restart capability, include the **disable** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for PIM Sparse Mode

PIM sparse mode continues to forward existing multicast packet streams during a graceful restart, but does not forward new streams until after the restart is complete. After a restart, the routing platform updates the forwarding state with any updates that were received from neighbors and occurred during the restart period. For example, the routing platform relearns the join and prune states of neighbors during the restart, but does not apply the changes to the forwarding table until after the restart.

PIM sparse mode-enabled routing platforms generate a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the IETF Internet draft *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*. When a routing platform receives PIM hellos containing generation identifiers on a point-to-point interface, Junos OS activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a PIM sparse mode-enabled routing platform restarts, it creates a new generation identifier and sends it to its neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires.

If a routing platform does not support generation identifiers or if PIM is enabled on multipoint interfaces, the PIM sparse mode graceful restart algorithm does not activate, and a default restart timer is used as the restart mechanism.

To configure the duration of the PIM graceful restart period, include the **restart-duration** statement at the **[edit protocols pim graceful-restart]** hierarchy level:

```
[edit]
protocols {
```

```
pim {  
  graceful-restart {  
    disable;  
    restart-duration seconds;  
  }  
}  
routing-options {  
  graceful-restart;  
}
```

To disable PIM sparse mode graceful restart capability, include the **disable** statement at the **[edit protocols pim graceful-restart]** hierarchy level.



NOTE: Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast reverse-path-forwarding (RPF) checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the **[edit protocols *protocol* traceoptions flag]** hierarchy level:

```
[edit protocols]  
isis {  
  traceoptions {  
    flag graceful-restart;  
  }  
}  
(ospf | ospf3) {  
  traceoptions {  
    flag graceful-restart;  
  }  
}
```

Related Documentation

- [Graceful Restart Concepts on page 3](#)
- [Graceful Restart System Requirements](#)
- [Graceful Restart and Routing Protocols](#)
- [Verifying Graceful Restart Operation on page 135](#)
- [Configuring Graceful Restart](#)

PART 2

Graceful Routing Engine Switchover

- [Using GRES on page 13](#)

CHAPTER 2

Using GRES

- [Understanding Graceful Routing Engine Switchover on page 13](#)
- [Graceful Routing Engine Switchover System Requirements on page 19](#)
- [Configuring Graceful Routing Engine Switchover on page 23](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 25](#)
- [Resetting Local Statistics on page 26](#)

Understanding Graceful Routing Engine Switchover

This topic contains the following sections:

- [Graceful Routing Engine Switchover Concepts on page 13](#)
- [Effects of a Routing Engine Switchover on page 18](#)

Graceful Routing Engine Switchover Concepts

The graceful Routing Engine switchover (GRES) feature in Junos OS enables a routing platform with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. GRES preserves interface and kernel information. Traffic is not interrupted. However, GRES does not preserve the control plane.



NOTE: On T Series routers, TX Matrix routers, and TX Matrix Plus routers, the control plane is preserved in case of GRES with nonstop active routing (NSR), and nearly 75 percent of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES.

Neighboring routers detect that the router has experienced a restart and react to the event in a manner prescribed by individual routing protocol specifications.

To preserve routing during a switchover, GRES must be combined with either:

- Graceful restart protocol extensions
- Nonstop active routing

Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur.



NOTE: Due to its synchronization requirements and logic, NSR/GRES performance is limited by the slowest Routing Engine in the system.

Mastership switches to the backup Routing Engine if:

- The master Routing Engine kernel stops operating.
- The master Routing Engine experiences a hardware failure.
- The administrator initiates a manual switchover.



NOTE: To quickly restore or to preserve routing protocol state information during a switchover, GRES must be combined with either graceful restart or nonstop active routing, respectively. For more information about graceful restart, see [“Graceful Restart Concepts” on page 3](#). For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 29](#).

If the backup Routing Engine does not receive a keepalive from the master Routing Engine after 2 seconds (4 seconds on M20 routers), it determines that the master Routing Engine has failed and: takes mastership.

The Packet Forwarding Engine:

- Seamlessly disconnects from the old master Routing Engine
- Reconnects to the new master Routing Engine
- Does not reboot
- Does not interrupt traffic

The new master Routing Engine and the Packet Forwarding Engine then become synchronized. If the new master Routing Engine detects that the Packet Forwarding Engine state is not up to date, it resends state update messages.



NOTE: If adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the *hold-time* for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.



NOTE: Successive Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to **Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset**. Do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the FPCs should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.



NOTE: Starting from Junos OS Release 14.2, when you perform GRES on MX Series routers, you must execute the **clear synchronous-ethernet wait-to-restore** operational mode command on the new master Routing Engine to clear the wait-to-restore timer on it. This is because the **clear synchronous-ethernet wait-to-restore** operational mode command clears the wait-to-restore timer only on the local Routing Engine.



NOTE: In a routing matrix with TX Matrix Plus router with 3D SIBs, for successive Routing Engine switchover, events must be a minimum of 900 seconds (15 minutes) apart after both Routing Engines have come up.

GRES must be performed on one line-card chassis (LCC) (of a TX Matrix router with 3D SIBs) at a time to avoid synchronization issues.

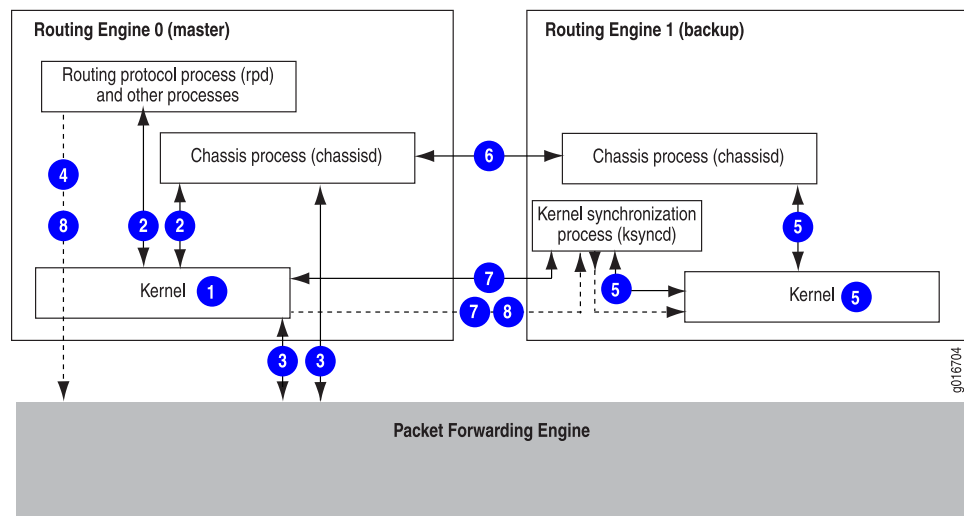


NOTE:

- We do *not* recommend performing a commit operation on the backup Routing Engine when GRES is enabled on the router or switch.
- We do *not* recommend enabling GRES on the backup Routing Engine in *any* scenario.

Figure 1 on page 16 shows the system architecture of graceful Routing Engine switchover and the process a routing platform follows to prepare for a switchover.

Figure 1: Preparing for a Graceful Routing Engine Switchover



NOTE: Check GRES readiness by executing both:

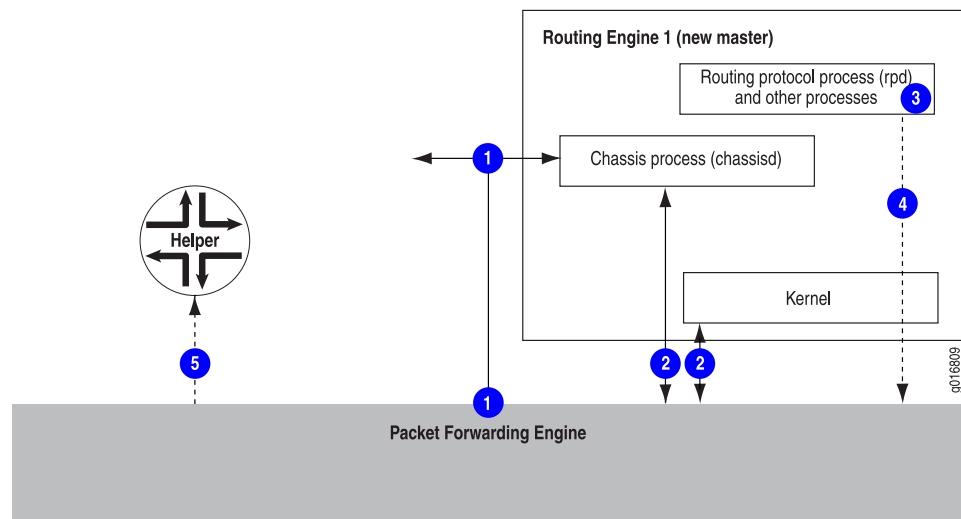
- The `request chassis routing-engine master switch check` command from the master Routing Engine
- The `show system switchover` command from the Backup Routing Engine

The switchover preparation process for GRES is as follows:

1. The master Routing Engine starts.
2. The routing platform processes (such as the chassis process [chassisd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts.
6. The system determines whether GRES has been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. After ksyncd completes the synchronization, all state information and the forwarding table are updated.

Figure 2 on page 17 shows the effects of a switchover on the routing (or switching) platform.

Figure 2: Graceful Routing Engine Switchover Process



When a switchover occurs, the switchover process is as follows:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master.
3. Routing platform processes that are not part of GRES (such as the routing protocol process [rpd]) restart.
4. State information learned from the point of the switchover is updated in the system.
5. If configured, graceful restart protocol extensions collect and restore routing information from neighboring peer *helper* routers.



NOTE: On T Series and M320 routers during GRES, the Switch Interface Boards (SIBs) are taken offline and restarted one by one. This is done to provide the Switch Processor Mezzanine Board (SPMB) that manages the SIB enough time to populate state information for its associated SIB. However, on a fully populated chassis where all FPCs are sending traffic at full line rate, there might be momentary packet loss during the switchover.



NOTE: When GRES is configured and the `restart chassis-control` command is executed on a TX Matrix Plus router with 3D SIBs, we cannot ascertain which Routing Engine becomes the master. This is because the `chassisd` process restarts with the execution of the `restart chassis-control` command. The `chassisd` process is responsible for maintaining and retaining mastership and when it is restarted, the new `chassisd` is processed based on the router or switch load. As a result, any one of the Routing Engines is made the master.

Effects of a Routing Engine Switchover

Table 3 on page 18 describes the effects of a Routing Engine switchover when different features are enabled:

- No high availability features
- Graceful Routing Engine switchover
- Graceful restart
- Nonstop active routing

Table 3: Effects of a Routing Engine Switchover

Feature	Benefits	Considerations
Dual Routing Engines only (no features enabled)	<ul style="list-style-type: none"> • When the switchover to the new master Routing Engine is complete, routing convergence takes place and traffic is resumed. 	<ul style="list-style-type: none"> • All physical interfaces are taken offline. • Packet Forwarding Engines restart. • The standby Routing Engine restarts the routing protocol process (rpd). • All hardware and interfaces are discovered by the new master Routing Engine. • The switchover takes several minutes. • All of the router's adjacencies are aware of the physical (interface alarms) and routing (topology) changes.
GRES enabled	<ul style="list-style-type: none"> • During the switchover, interface and kernel information is preserved. • The switchover is faster because the Packet Forwarding Engines are not restarted. 	<ul style="list-style-type: none"> • The new master Routing Engine restarts the routing protocol process (rpd). • All hardware and interfaces are acquired by a process that is similar to a warm restart. • All adjacencies are aware of the router's change in state.
GRES <i>and</i> nonstop active routing enabled	<ul style="list-style-type: none"> • Traffic is not interrupted during the switchover. • Interface and kernel information are preserved. 	<ul style="list-style-type: none"> • Unsupported protocols must be refreshed using the normal recovery mechanisms inherent in each protocol.

Table 3: Effects of a Routing Engine Switchover (*continued*)

Feature	Benefits	Considerations
GRES and graceful restart enabled	<ul style="list-style-type: none"> Traffic is not interrupted during the switchover. Interface and kernel information are preserved. Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers. 	<ul style="list-style-type: none"> Neighbors are required to support graceful restart, and a wait interval is required. The routing protocol process (rpd) restarts. For certain protocols, a significant change in the network can cause graceful restart to stop. If adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart can stop and cause interruptions in traffic.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Graceful Routing Engine Switchover System Requirements on page 19](#)
- [Configuring Graceful Routing Engine Switchover on page 23](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 25](#)
- [Requirements for Routers with a Backup Router Configuration](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)
- [hold-time](#)

Graceful Routing Engine Switchover System Requirements

Graceful Routing Engine switchover is supported on all routing (or switching) platforms that contain dual Routing Engines. All Routing Engines configured for graceful Routing Engine switchover must run the same Junos OS release. Hardware and software support for graceful Routing Engine switchover is described in the following sections:

- [Graceful Routing Engine Switchover Platform Support on page 19](#)
- [Graceful Routing Engine Switchover Feature Support on page 20](#)
- [Graceful Routing Engine Switchover DPC Support on page 22](#)
- [Graceful Routing Engine Switchover and Subscriber Access on page 22](#)
- [Graceful Routing Engine Switchover PIC Support on page 22](#)

Graceful Routing Engine Switchover Platform Support

To enable graceful Routing Engine switchover, your system must meet these minimum requirements:

- M20 and M40e routers—Junos OS Release 5.7 or later
- M10i router—Junos OS Release 6.1 or later
- M320 router—Junos OS Release 6.2 or later
- T320 router, T640 router, and TX Matrix router—Junos OS Release 7.0 or later
- M120 router—Junos OS Release 8.2 or later
- MX960 router—Junos OS Release 8.3 or later
- MX480 router—Junos OS Release 8.4 or later (8.4R2 recommended)
- MX240 router—Junos OS Release 9.0 or later
- PTX5000 router—Junos OS Release 12.1X48 or later
- Standalone T1600 router—Junos OS Release 8.5 or later
- Standalone T4000 router—Junos OS Release 12.1R2 or later
- TX Matrix Plus router—Junos OS Release 9.6 or later
- TX Matrix Plus router with 3D SIBs—Junos Release 13.1 or later
- EX Series switches with dual Routing Engines or in a Virtual Chassis — Junos OS Release 9.2 or later for EX Series switches
- QFX Series switches in a Virtual Chassis —Junos OS Release 13.2 or later for the QFX Series
- EX Series or QFX Series switches in a Virtual Chassis Fabric —Junos OS Release 13.2X51-D20 or later for the EX Series and QFX Series switches

For more information about support for graceful Routing Engine switchover, see the sections that follow.

Graceful Routing Engine Switchover Feature Support

Graceful Routing Engine switchover supports most Junos OS features in Release 5.7 and later. Particular Junos OS features require specific versions of Junos OS. See [Table 4 on page 20](#).

Table 4: Graceful Routing Engine Switchover Feature Support

Application	Junos OS Release
Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP) and aggregated SONET interfaces	6.2
Asynchronous Transfer Mode (ATM) virtual circuits (VCs)	6.2
Logical systems	6.3
NOTE: In Junos OS Release 9.3 and later, the logical router feature is renamed to logical system.	
Multicast	6.4 (7.0 for TX Matrix router)

Table 4: Graceful Routing Engine Switchover Feature Support (*continued*)

Application	Junos OS Release
Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR)	7.0
Automatic Protection Switching (APS)—The current active interface (either the designated working or the designated protect interface) remains the active interface during a Routing Engine switchover.	7.4
Point-to-multipoint Multiprotocol Label Switching MPLS LSPs (transit only)	7.4
Compressed Real-Time Transport Protocol (CRTP)	7.6
Virtual private LAN service (VPLS)	8.2
Ethernet Operation, Administration, and Management (OAM) as defined by IEEE 802.3ah	8.5
Extended DHCP relay agent	8.5
Ethernet OAM as defined by IEEE 802.1ag	9.0
Packet Gateway Control Protocol (PGCP) process (pgcpd) on Multiservices 500 PICs on T640 routers.	9.0
Subscriber access	9.4
Layer 2 Circuit and LDP-based VPLS pseudowire redundant configuration	9.6

The following constraints apply to graceful Routing Engine switchover feature support:

- When graceful Routing Engine switchover and aggregated Ethernet interfaces are configured in the same system, the aggregated Ethernet interfaces must not be configured for fast-polling LACP. When fast polling is configured, the LACP polls time out at the remote end during the Routing Engine mastership switchover. When LACP polling times out, the aggregated link and interface are disabled. The Routing Engine mastership change is fast enough that standard and slow LACP polling do not time out during the procedure. However, note that this restriction does not apply to MX Series Routers that are running Junos OS Release 9.4 or later and have distributed periodic packet management (PPM) enabled—which is the default configuration—on them. In such cases, you can configure graceful Routing Engine switchover and have aggregated Ethernet interfaces configured for fast-polling LACP on the same device.
- When a graceful Routing Engine switchover occurs, the VRRP state does not change. VRRP is supported by graceful Routing Engine switchover only in the case that PPM delegation is enabled (which the default).

Graceful Routing Engine Switchover DPC Support

Graceful Routing Engine switchover supports all Dense Port Concentrators (DPCs) on the MX Series 3D Universal Edge Routers running the appropriate version of Junos OS as shown in “[Graceful Routing Engine Switchover Platform Support](#)” on page 19. For more information about DPCs, see the *MX Series DPC Guide*.

Graceful Routing Engine Switchover and Subscriber Access

Graceful Routing Engine switchover currently supports most of the features directly associated with dynamic DHCP and dynamic PPPoE subscriber access. Graceful Routing Engine switchover also supports the unified in-service software upgrade (ISSU) for the DHCP access model and the PPPoE access model used by subscriber access.

Graceful Routing Engine Switchover PIC Support

Graceful Routing Engine switchover is supported on most PICs, except for the services PICs listed in this section. The PIC must be on a supported routing platform running the appropriate version of Junos OS. For information about FPC types, FPC/PIC compatibility, and the initial Junos OS Release in which an FPC supported a particular PIC, see the PIC guide for your router platform.

The following constraints apply to graceful Routing Engine switchover support for services PICs:

- You can include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with Adaptive Services, Multiservices, and Tunnel Services PICs configured on it and successfully commit the configuration. However, all services on these PICs—except the Layer 2 service packages and extension-provider and SDK applications on Multiservices PICs—are reset during a switchover.
- Graceful Routing Engine switchover is not supported on any Monitoring Services PICs or Multilink Services PICs. If you include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with either of these PIC types configured on it and issue the **commit** command, the commit fails.
- Graceful Routing Engine switchover is not supported on Multiservices 400 PICs configured for monitoring services applications. If you include the **graceful-switchover** statement, the commit fails.



NOTE: When an unsupported PIC is online, you cannot enable graceful Routing Engine switchover. If graceful Routing Engine switchover is already enabled, an unsupported PIC cannot come online.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Understanding Graceful Routing Engine Switchover on page 13](#)
- [Configuring Graceful Routing Engine Switchover on page 23](#)

- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 25](#)
- [Requirements for Routers with a Backup Router Configuration](#)

Configuring Graceful Routing Engine Switchover

This section contains the following topics:

- [Enabling Graceful Routing Engine Switchover on page 23](#)
- [Configuring Graceful Routing Engine Switchover with Graceful Restart on page 23](#)
- [Synchronizing the Routing Engine Configuration on page 23](#)
- [Verifying Graceful Routing Engine Switchover Operation on page 25](#)

Enabling Graceful Routing Engine Switchover

By default, graceful Routing Engine switchover (GRES) is disabled. To configure GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

```
[edit chassis redundancy]
graceful-switchover;
```

When you enable GRES, the command-line interface (CLI) indicates which Routing Engine you are using. For example:

```
{master} [edit]
user@host#
```

To disable GRES, delete the **graceful-switchover** statement from the **[edit chassis redundancy]** hierarchy level.

Configuring Graceful Routing Engine Switchover with Graceful Restart

When using GRES with Graceful Restart, if adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the *hold-time* for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.

Synchronizing the Routing Engine Configuration



NOTE: A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure GRES, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.

Only when you enable the graceful Routing Engine switchover, you can copy the running Junos OS version of the master Routing Engine to the backup Routing Engine.



NOTE: If the system is in ISSU state, you cannot copy the running Junos OS version of the master Router Engine.

You can enable automatic synchronization of the master Routing Engine configuration with the backup Routing Engine by including the events CHASSISD_SNMP_TRAP7 statement at the [edit event-options policy *policy-name*] hierarchy level.

CHASSISD_SNMP_TRAP7 is a system event logging message that the chassis process (chassisd) generates a Simple Network Management Protocol (SNMP) trap with the seven indicated argument-value pairs. An example of an event script to trigger automatic synchronization of master to the backup Routing Engine is as follows:

```
[edit event-options]
policy UPGRADE-BACKUPRE {
  events CHASSISD_SNMP_TRAP7;
  attributes-match {
    CHASSISD_SNMP_TRAP7.value5 matches "Routing Engine";
    CHASSISD_SNMP_TRAP7.trap matches "Fru Online";
    CHASSISD_SNMP_TRAP7.argument5 matches jnxFruName;
  }
  then {
    event-script auto-image-upgrade.slax {
      arguments {
        trap "${$.trap}";
        value5 "${$.value5}";
        argument5 "${$.argument5}";
      }
    }
  }
  event-script {
    file auto-image-upgrade.slax;
  }
}
```

After receiving this event, the event policy on the master Router Engine is triggered and the image available in the */var/sw/pkg* path is pushed to the backup Router Engine upgrade. During script execution, the image is copied to the backup Routing Engine's */var/sw/pkg* path.



NOTE: If the image is not available in the */var/sw/pkg* path, the script is terminated with an appropriate syslog message.

If the Routing Engine is running at the Junos OS Release 13.2 or later, the Junos automation scripts is synchronized automatically.

After the master Router Engine is rebooted, the event script available at the */usr/libexec/scripts/event/auto-image-upgrade.slax* must be copied to the */var/db/scripts/event* path.

Verifying Graceful Routing Engine Switchover Operation

To verify whether GRES is enabled on the backup Routing Engine, issue the **show system switchover** command. When the output of the command indicates that the **Graceful switchover** field is set to **On**, GRES is operational. The status of the kernel database and configuration database synchronization between Routing Engines is also provided. For example:

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady state
```



NOTE: You must issue the **show system switchover** command on the backup Routing Engine. This command is not supported on the master Routing Engine.

For more information about the **show system switchover** command, see the [CLI Explorer](#).

Related Documentation

- [Understanding Graceful Routing Engine Switchover on page 13](#)
- [Graceful Routing Engine Switchover System Requirements on page 19](#)
- [Requirements for Routers with a Backup Router Configuration](#)
- [Resetting Local Statistics on page 26](#)
- [graceful-switchover](#)
- [graceful-switchover on page 101](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)
- [hold-time](#)

Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)

In a Virtual Chassis, one member switch is assigned the master role and has the master Routing Engine. Another member switch is assigned the backup role and has the backup Routing Engine. Graceful Routing Engine switchover (GRES) enables the master and backup Routing Engines in a Virtual Chassis configuration to switch from the master to backup without interruption to packet forwarding. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and the forwarding state.

To set up the Virtual Chassis configuration to use graceful Routing Engine switchover (GRES):

1. Set up a minimum of two switches in a Virtual Chassis configuration with mastership priority of 255:

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

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

2. Set up graceful Routing Engine switchover:

```
[edit]
```

```
user@switch# set chassis redundancy graceful-switchover
```

Commit the configuration.



NOTE: We recommend that you use the `commit synchronize` command to save any configuration changes that you make to a multimember Virtual Chassis.

**Related
Documentation**

- *Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet*
- *High Availability Features for EX Series Switches Overview*
- *Understanding EX Series Virtual Chassis Configuration*
- *Understanding QFX Series Virtual Chassis*

Resetting Local Statistics

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, accounting information, and trace options information are not replicated to the backup Routing Engine.

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. 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. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the **clear** command to clear statistics and protocol database information, see the [CLI Explorer](#).



NOTE: The `clear firewall` command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover.

**Related
Documentation**

- [Understanding Graceful Routing Engine Switchover on page 13](#)
- [Configuring Graceful Routing Engine Switchover on page 23](#)

PART 3

Nonstop Bridging and Routing

- [Using Nonstop Bridging and Routing on page 29](#)

CHAPTER 3

Using Nonstop Bridging and Routing

- [Nonstop Active Routing Concepts on page 29](#)
- [Nonstop Active Routing System Requirements on page 32](#)
- [Nonstop Bridging Concepts on page 43](#)
- [Nonstop Bridging System Requirements on page 45](#)
- [Configuring Nonstop Active Routing on Switches on page 46](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 47](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 49](#)
- [Resetting Local Statistics on page 50](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 51](#)

Nonstop Active Routing Concepts

Nonstop active routing (NSR) uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, NSR also saves routing protocol information by running the routing protocol process (rpd) on the backup Routing Engine. By saving this additional information, NSR is self-contained and does not rely on helper routers (or switches) to assist the routing platform in restoring routing protocol information. NSR is advantageous in networks in which neighbor routers (or switches) do not support graceful restart protocol extensions. As a result of this enhanced functionality, NSR is a natural replacement for graceful restart.

If you have NSR configured, it is never valid to issue the **restart routing** command in any form on the NSR master Routing Engine. Doing so results in a loss of protocol adjacencies and neighbors and a drop in traffic.




NOTE: To use NSR, you must first enable GRES on your routing (or switching) platform. For more information about GRES, see [“Understanding Graceful Routing Engine Switchover” on page 13](#).



NOTE: Due to its synchronization requirements and logic, NSR/GRES performance is limited by the slowest Routing Engine in the system.

Figure 3 on page 30 shows the system architecture of nonstop active routing and the process a routing (or switching) platform follows to prepare for a switchover.

Figure 3: Nonstop Active Routing Switchover Preparation Process




The switchover preparation process for NSR follows these steps:

1. The master Routing Engine starts.
2. The routing (or switching) platform processes on the master Routing Engine (such as the chassis process [chassisd] and the routing protocol process [rpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the routing protocol process (rpd).
6. The system determines whether graceful Routing Engine switchover and nonstop active routing have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the routing protocol processes on the master and backup Routing Engines.

Figure 4 on page 31 shows the effects of a switchover on the routing platform.

Figure 4: Nonstop Active Routing During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the routing protocol process (rpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and routing are continued during the switchover, resulting in minimal packet loss.
4. Peer routers (or switches) continue to interact with the routing platform as if no change had occurred. Routing adjacencies and session state relying on underlying routing information are preserved and not reset.



CAUTION: We recommend that you do not restart the routing protocol process (rpd) on master Routing Engine after enabling NSR, as it disrupts the protocol adjacency/peering sessions, resulting in traffic loss.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Active Routing System Requirements on page 32](#)
- [Configuring Nonstop Active Routing](#)

- [Configuring Nonstop Active Routing on Switches on page 46](#)

Nonstop Active Routing System Requirements

This section contains the following topics:

- [Nonstop Active Routing Platform and Switching Platform Support on page 32](#)
- [Nonstop Active Routing Protocol and Feature Support on page 33](#)
- [Nonstop Active Routing BFD Support on page 36](#)
- [Nonstop Active Routing BGP Support on page 37](#)
- [Nonstop Active Routing Layer 2 Circuit and VPLS Support on page 38](#)
- [Nonstop Active Routing PIM Support on page 38](#)
- [Nonstop Active Routing MSDP Support on page 41](#)
- [Nonstop Active Routing Support for RSVP-TE LSPs on page 41](#)

Nonstop Active Routing Platform and Switching Platform Support

Table 5 on page 32 lists the platforms that support nonstop active routing (NSR).

Table 5: Nonstop Active Routing Platform Support

Platform	Junos OS Release
M10i router	8.4 or later
M20 router	8.4 or later
M40e router	8.4 or later
M120 router	9.0 or later
M320 router	8.4 or later
MX Series routers	9.0 or later
PTX Series Packet Transport Routers	12.1R4 or later

NOTE:

Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:

- Labeled BGP
- Layer 2 VPNs excluding Layer 2 interworking (Layer 2 switching)
- Layer 3 VPNs
- LDP
- RSVP

Table 5: Nonstop Active Routing Platform Support (*continued*)

Platform	Junos OS Release
PTX Series Packet Transport Routers	12.1R4 or later
<p>NOTE:</p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> • Labeled BGP • Layer 2 VPNs excluding Layer 2 interworking (Layer 2 stitching) • Layer 3 VPNs • LDP • RSVP 	
PTX Series Packet Transport Routers	12.1R4 or later
<p>NOTE:</p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> • Labeled BGP • Layer 2 VPNs excluding Layer 2 interworking (Layer 2 stitching) • Layer 3 VPNs • LDP • RSVP 	
T320 router, T640 router, and TX Matrix router	8.4 or later
Standalone T1600 router	8.5 or later
Standalone T4000 router	12.1R2 or later
TX Plus Matrix router	10.0 or later
TX Plus Matrix router with 3D SIBs	13.1 or later
EX Series switch with dual Routing Engines or in a Virtual Chassis	10.4 or later for EX Series switches
EX Series or QFX Series switches in a Virtual Chassis Fabric	13.2X51-D20 or later for the EX Series and QFX Series switches



NOTE: All Routing Engines configured for nonstop active routing must be running the same Junos OS release.

Nonstop Active Routing Protocol and Feature Support

Table 6 on page 34 lists the protocols that are supported by nonstop active routing.

Table 6: Nonstop Active Routing Protocol and Feature Support

Protocol	Junos OS Release
Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP)	9.4 or later
Bidirectional Forwarding Detection (BFD) For more information, see “Nonstop Active Routing BFD Support” on page 36.	8.5 or later
BGP For more information, see “Nonstop Active Routing BGP Support” on page 37.	8.4 or later
Labeled BGP (PTX Series Packet Transport Routers: only)	12.1R4 or later
IS-IS	8.4 or later
LDP	8.4 or later
LDP-based virtual private LAN service (VPLS)	9.3 or later
LDP OAM (operation, administration, and management) features	9.6 or later
LDP (PTX Series Packet Transport Routers only) Nonstop active routing support for LDP includes: <ul style="list-style-type: none"> • LDP unicast transit LSPs • LDP egress LSPs for labeled internal BGP (IBGP) and external BGP (EBGP) • LDP over RSVP transit LSPs • LDP transit LSPs with indexed next hops • LDP transit LSPs with unequal cost load balancing NOTE: Nonstop active routing is not supported for LDP Point-to-Multipoint LSPs and LDP ingress LSPs.	12.3R4 or later
Layer 2 circuits	(on LDP-based VPLS) 9.2 or later (on RSVP-TE LSP) 11.1 or later
Layer 2 VPNs	9.1 or later
Layer 2 VPNs (PTX Series Packet Transport Routers only)	12.1R4 or later
NOTE: Nonstop active routing is not supported for Layer 2 interworking (Layer 2 stitching).	

Table 6: Nonstop Active Routing Protocol and Feature Support (*continued*)

Protocol	Junos OS Release
Layer 3 VPNs (see the first Note after this table for restrictions)	9.2 or later
Nonstop active routing support for Layer 3 VPNs include:	
<ul style="list-style-type: none"> • IPv4 labeled-unicast (ingress or egress) • IPv4-vpn unicast (ingress or egress) • IPv6 labeled-unicast (ingress or egress) • IPv6-vpn unicast (ingress or egress) 	
Layer 3 VPNs (PTX Series Packet Transport Routers only)	12.1R4 or later
Multicast Source Discovery Protocol (MSDP)	12.1 or later
For more information, see “Nonstop Active Routing MSDP Support” on page 41 .	
OSPF/OSPFv3	8.4 or later
Protocol Independent Multicast (PIM)	(for IPv4) 9.3 or later
For more information, see “Nonstop Active Routing PIM Support” on page 38 .	(for IPv6) 10.4 or later
RIP and RIP next generation (RIPng)	9.0 or later
RSVP (PTX Series Packet Transport Routers only)	12.1R4 or later
Nonstop active routing support for RSVP includes:	
<ul style="list-style-type: none"> • Point-to-Multipoint LSPs <ul style="list-style-type: none"> • RSVP Point-to-Multipoint ingress, transit, and egress LSPs using existing non-chained next hop. • RSVP Point-to-Multipoint transit LSPs using composite next hops for Point-to-Multipoint label routes. • Point-to-Point LSPs <ul style="list-style-type: none"> • RSVP Point-to-Point ingress, transit, and egress LSPs using non-chained next hops. • RSVP Point-to-Point transit LSPs using chained composite next hops. 	
RSVP-TE LSP	9.5 or later
For more information, see “Nonstop Active Routing Support for RSVP-TE LSPs” on page 41 .	
VPLS	(LDP-based) 9.1 or later (RSVP-TE-based) 11.2 or later
VRRP	13.2 or later

Table 6: Nonstop Active Routing Protocol and Feature Support (*continued*)

Protocol	Junos OS Release
VRRP	13.2 or later



NOTE: Layer 3 VPN support does not include dynamic GRE tunnels, multicast VPNs, or BGP flow routes.



NOTE: If you configure a protocol that is not supported by nonstop active routing, the protocol operates as usual. When a switchover occurs, the state information for the unsupported protocol is not preserved and must be refreshed using the normal recovery mechanisms inherent in the protocol.



NOTE: On routers that have logical systems configured on them, only the master logical system supports nonstop active routing.



NOTE: On EX9214 switches, the VRRP master state might change during graceful Routing Engine switchover, even when nonstop active routing is enabled.

Nonstop Active Routing BFD Support

Nonstop active routing supports the Bidirectional Forwarding Detection (BFD) protocol, which uses the topology discovered by routing protocols to monitor neighbors. The BFD protocol is a simple hello mechanism that detects failures in a network. Because BFD is streamlined to be efficient at fast liveness detection, when it is used in conjunction with routing protocols, routing recovery times are improved. With nonstop active routing enabled, BFD session states are not restarted when a Routing Engine switchover occurs.



NOTE: BFD session states are saved only for clients using aggregate or static routes or for BGP, IS-IS, OSPF/OSPFv3, or PIM.

When a BFD session is distributed to the Packet Forwarding Engine, BFD packets continue to be sent during a Routing Engine switchover. If nondistributed BFD sessions are to be kept alive during a switchover, you must ensure that the session failure detection time is greater than the Routing Engine switchover time. The following BFD sessions are not distributed to the Packet Forwarding Engine: multihop sessions, tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping. The minimum-interval configuration statement is a BFD liveness detection parameter.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions, and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, specify a minimum interval of 10 seconds for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

Nonstop Active Routing BGP Support

Nonstop active routing BGP support is subject to the following conditions:

- You must include the **path-selection external-router-ID** statement at the **[edit protocols bgp]** hierarchy level to ensure consistent path selection between the master and backup Routing Engines during and after the nonstop active routing switchover.
- You must include the **advertise-from-main-vpn-tables** statement at the **[edit protocols bgp]** hierarchy level to prevent BGP sessions from going down when route reflector (RR) or autonomous system border router (ASBR) functionality is enabled or disabled on a routing device that has VPN address families configured.
- BGP session uptime and downtime statistics are not synchronized between the primary and backup Routing Engines during Nonstop Active Routing and ISSU. The backup Routing Engine maintains its own session uptime based on the time when the backup first becomes aware of the established sessions. For example, if the backup Routing Engine is rebooted (or if you run **restart routing** on the backup Routing Engine), the backup's uptime is a short duration, because the backup has just learned about the established sessions. If the backup is operating when the BGP sessions first come up on the primary, the uptime on the primary and the uptime on the backup are almost the same duration. After a Routing Engine switchover, the new master continues from the time left on the standby Routing Engine.
- If the BGP peer in the master Routing Engine has negotiated address-family capabilities that are not supported for nonstop active routing, then the corresponding BGP neighbor state on the backup Routing Engine shows as idle. On switchover, the BGP session is reestablished from the new master Routing Engine.

Only the following address families are supported for nonstop active routing.



NOTE: Address families are supported only on the main instance of BGP. Only unicast is supported on VRF instances.

- inet unicast
 - inet labeled-unicast
 - inet multicast
 - inet6 labeled-unicast
 - inet6 multicast
 - inet6 unicast
 - route-target
 - l2vpn signaling
 - inet6-vpn unicast
 - inet-vpn unicast
 - inet-mdt
 - iso-vpn
- BGP route dampening does not work on the backup Routing Engine when nonstop active routing is enabled.

Nonstop Active Routing Layer 2 Circuit and VPLS Support

Nonstop active routing supports Layer 2 circuit and VPLS on both LDP-based and RSVP-TE-based networks. Nonstop active routing support enables the backup Routing Engine to track the label advertised by Layer 2 circuit and VPLS on the primary Routing Engine, and to use the same label after the Routing Engine switchover.

in Junos OS Release 9.6 and later, nonstop active routing support is extended to the Layer 2 circuit and LDP-based VPLS pseudowire redundant configurations.

Nonstop Active Routing PIM Support

Nonstop active routing supports Protocol Independent Multicast (PIM) with stateful replication on backup Routing Engines. State information replicated on the backup Routing Engine includes information about neighbor relationships, join and prune events, rendezvous point (RP) sets, synchronization between routes and next hops, multicast session states, and the forwarding state between the two Routing Engines.



NOTE: Nonstop active routing for PIM is supported for IPv4 on Junos OS Release 9.3 and later, and for IPv6 on Junos OS Release 10.4 and later. Starting with Release 11.1, Junos OS also supports nonstop active routing for PIM on devices that have both IPv4 and IPv6 configured on them.

To configure nonstop active routing for PIM, include the same statements in the configuration as for other protocols: the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. To trace PIM nonstop active routing events, include the **flag nsr-synchronization** statement at the **[edit protocols pim traceoptions]** hierarchy level.



NOTE: The **clear pim join**, **clear pim register**, and **clear pim statistics** operational mode commands are not supported on the backup Routing Engine when nonstop active routing is enabled.

Nonstop active routing support varies for different PIM features. The features fall into the following three categories: supported features, unsupported features, and incompatible features.

Supported features:

- Auto-RP



NOTE: Nonstop active routing PIM support on IPv6 does not support auto-RP because IPv6 does not support auto-RP.

- Bootstrap router (BSR)
- Static RPs
- Embedded RP on non-RP IPv6 routers
- Local RP



NOTE: RP set information synchronization is supported for local RP and BSR (on IPv4 and IPv6), autoRP (on IPv4), and embedded RP (on IPv6).

- BFD
- Dense mode
- Sparse mode
- Source-specific multicast (SSM)
- Draft Rosen multicast VPNs (MVPNs)

- Anycast RP (anycast RP set information synchronization and anycast RP register state synchronization on IPv4 and IPv6 configurations)
- Flow maps
- Unified ISSU
- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies
- Upstream assert synchronization
- PIM join load balancing

Starting with Release 12.2, Junos OS extends the nonstop active routing PIM support to draft Rosen MVPNs. Nonstop active routing PIM support for draft Rosen MVPNs enables nonstop active routing-enabled devices to preserve draft Rosen MPVN-related information—such as default and data multicast distribution tree (MDT) states—across switchovers. In releases earlier than 12.2, nonstop active routing PIM configuration was incompatible with draft Rosen MPVN configuration.

The backup Routing Engine sets up the default MDT based on the configuration and the information it receives from the master Routing Engine, and keeps updating the default MDT state information.

However, for data MDTs, the backup Routing Engine relies on the master Routing Engine to provide updates when data MDTs are created, updated, or deleted. The backup Routing Engine neither monitors data MDT flow rates nor triggers a data MDT switchover based on variations in flow rates. Similarly, the backup Routing Engine does not maintain the data MDT delay timer or timeout timer. It does not send MDT join TLV packets for the data MDTs until it takes over as the master Routing Engine. After the switchover, the new master Routing Engine starts sending MDT join TLV packets for each data MDT, and also resets the data MDT timers. Note that the expiration time for the timers might vary from the original values on the previous master Routing Engine.

Starting with Release 12.3, Junos OS extends the Protocol Independent Multicast (PIM) nonstop active routing support to IGMP-only interfaces.

In Junos OS releases earlier than 12.3, the PIM joins created on IGMP-only interfaces were not replicated on the backup Routing Engine. Thus, the corresponding multicast routes were marked as pruned (meaning discarded) on the backup Routing Engine. Because of this limitation, after a switchover, the new master Routing Engine had to wait for the IGMP module to come up and start receiving reports to create PIM joins and to install multicast routes. This caused traffic loss until the multicast joins and routes were reinstated.

However, in Junos OS Release 12.3 and later, the multicast joins on the IGMP-only interfaces are mapped to PIM states, and these states are replicated on the backup Routing Engine. If the corresponding PIM states are available on the backup, the multicast routes are marked as forwarding on the backup Routing Engine. This enables uninterrupted traffic flow after a switchover. This enhancement covers IGMPv2, IGMPv3, MLDv1, and MLDv2 reports and leaves.

Unsupported features: You can configure the following PIM features on a router along with nonstop active routing, but they function as if nonstop active routing is not enabled. In other words, during Routing Engine switchover and other outages, their state information is not preserved, and traffic loss is to be expected.

- Internet Group Management Protocol (IGMP) exclude mode
- IGMP snooping

Nonstop active routing is not supported for next-generation MVPNs with PIM provider tunnels. The commit operation fails if the configuration includes both nonstop active routing and next-generation MVPNs with PIM provider tunnels.

Junos OS provides a configuration statement that disables nonstop active routing for PIM only, so that you can activate incompatible PIM features and continue to use nonstop active routing for the other protocols on the router. Before activating an incompatible PIM feature, include the **nonstop-routing disable** statement at the **[edit protocols pim]** hierarchy level. Note that in this case, nonstop active routing is disabled for all PIM features, not just incompatible features.

Nonstop Active Routing MSDP Support

Starting with Release 12.1, Junos OS extends nonstop active routing support to the Multicast Source Discovery Protocol (MSDP).

Nonstop active routing support for MSDP preserves the following MSDP-related information across the switchover:

- MSDP configuration and peer information
- MSDP peer socket information
- Source-active and related information

However, note that the following restrictions or limitations apply to nonstop active routing MSDP support:

- Because the backup Routing Engine learns the active source information by processing the source-active messages from the network, synchronizing of source active information between the master and backup Routing Engines might take up to 60 seconds. So, no planned switchover is allowed within 60 seconds of the initial replication of the sockets.
- Similarly, Junos OS does not support two planned switchovers within 240 seconds of each other.

Junos OS enables you to trace MSDP nonstop active routing events by including the **flag nsr-synchronization** statement at the **[edit protocols msdp traceoptions]** hierarchy level.

Nonstop Active Routing Support for RSVP-TE LSPs

Junos OS extends nonstop active routing support to label-switching routers (LSRs) and Layer 2 Circuits that are part of an RSVP-TE LSP. Nonstop active routing support on LSRs ensures that the master to backup Routing Engine switchover on an LSR remains

transparent to the network neighbors and that the LSP information remains unaltered during and after the switchover.

You can use the **show rsvp version** command to view the nonstop active routing mode and state on an LSR. Similarly, you can use the **show mpls lsp** and **show rsvp session** commands on the standby Routing Engine to view the state recreated on the standby Routing Engine.

The Junos OS nonstop active routing feature is also supported on RSVP point-to-multipoint LSPs. Nonstop active routing support for RSVP point-to-multipoint egress and transit LSPs was added in Junos OS Release 11.4, and for ingress LSPs in Release 12.1. During the switchover, the LSP comes up on the backup Routing Engine that shares and synchronizes the state information with the master Routing Engine before and after the switchover. Nonstop active routing support for point-to-multipoint transit and egress LSPs ensures that the switchover remains transparent to the network neighbors, and preserves the LSP information across the switchover.

However, Junos OS nonstop active routing support for RSVP point-to-multipoint LSPs does not include support for dynamically created point-to-multipoint LSPs, such as VPLS.

Starting with Release 14.1, Junos OS extends nonstop active routing support to the next-generation multicast VPNs.

The **show rsvp session detail** command enables you to check the point-to-multipoint LSP remerge state information (**P2MP LSP re-merge**; possible values are **head**, **member**, and **none**).

However, Junos OS does not support nonstop active routing for the following features:

- Generalized Multiprotocol Label Switching (GMPLS) and LSP hierarchy
- Interdomain or loose-hop expansion LSPs
- BFD liveness detection

Nonstop active routing support for RSVP-TE LSPs is subject to the following limitations and restrictions:

- Detour LSPs are not maintained across a switchover and so, detour LSPs might fail to come back online after the switchover.
- Control plane statistics corresponding to the **show rsvp statistics** and **show rsvp interface detail | extensive** commands are not maintained across Routing Engine switchovers.
- Statistics from the backup Routing Engine are not reported for **show mpls lsp statistics** and **monitor mpls label-switched-path** commands. However, if a switchover occurs, the backup Routing Engine, after taking over as the master, starts reporting statistics. Note that the **clear statistics** command issued on the old master Routing Engine does not have any effect on the new master Routing Engine, which reports statistics, including any uncleared statistics.
- State timeouts might take additional time during nonstop active routing switchover. For example, if a switchover occurs after a neighbor has missed sending two hello

messages to the master, the new master Routing Engine waits for another three hello periods before timing out the neighbor.

- On the RSVP ingress router, if you configure auto-bandwidth functionality, the bandwidth adjustment timers are set in the new master after the switchover. This causes a one-time increase in the length of time required for the bandwidth adjustment after the switchover occurs.
- RSVP ingress LSPs that have BFD liveness detection enabled on them do not come up on the backup Routing Engine during the switchover. Such BFD-enabled LSPs have to be reestablished after the switchover.
- Backup LSPs —LSPs that are established between the point of local repair (PLR) and the merge point after a node or link failure—are not preserved during a Routing Engine switchover.
- When nonstop active routing is enabled, graceful restart is not supported. However, graceful restart helper mode is supported.

Related Documentation

- [Nonstop Active Routing Concepts on page 29](#)
- [Configuring Nonstop Active Routing](#)
- [Configuring Nonstop Active Routing on Switches on page 46](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 51](#)


Nonstop Bridging Concepts

Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop bridging also saves Layer 2 Control Protocol (L2CP) information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine.



NOTE: To use nonstop bridging, you must first enable graceful Routing Engine switchover on your routing (or switching) platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover” on page 13](#).

Figure 5 on page 44 shows the system architecture of nonstop bridging and the process a routing (or switching) platform follows to prepare for a switchover.

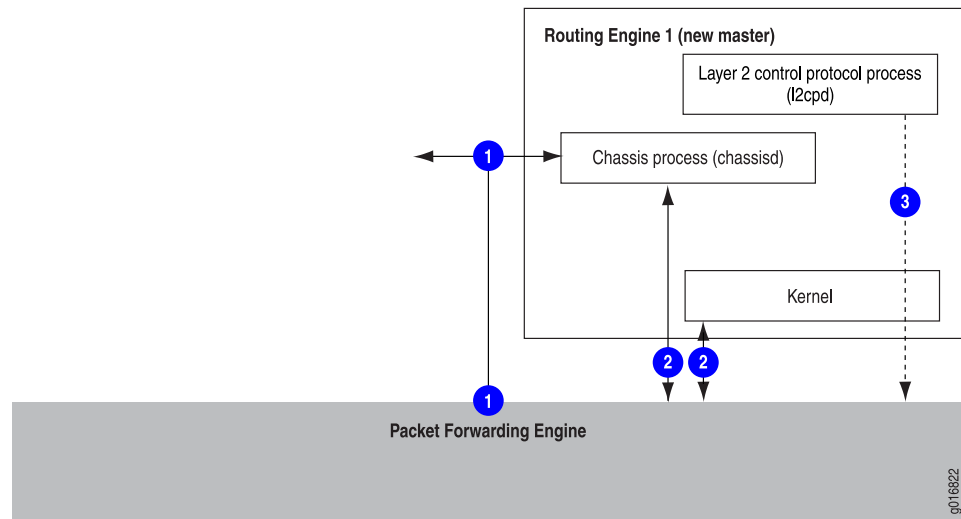
Figure 5: Nonstop Bridging Switchover Preparation Process

The switchover preparation process for nonstop bridging follows these steps:

1. The master Routing Engine starts.
2. The routing platform processes on the master Routing Engine (such as the chassis process [chassisd] and the Layer 2 Control Protocol process [l2cpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the Layer 2 Control Protocol process (l2cpd).
6. The system determines whether graceful Routing Engine switchover and nonstop bridging have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the l2cpds on the master and backup Routing Engines.

Figure 6 on page 45 shows the effects of a switchover on the routing platform.

Figure 6: Nonstop Bridging During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the Layer 2 Control Protocol process (l2cpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and bridging are continued during the switchover, resulting in minimal packet loss.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Bridging System Requirements on page 45](#)
- [Configuring Nonstop Bridging](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 49](#)

Nonstop Bridging System Requirements

This topic contains the following sections:

- [Platform Support on page 45](#)
- [Protocol Support on page 46](#)

Platform Support

Nonstop bridging is supported on MX Series 3D Universal Edge Routers. Your system must be running Junos OS Release 8.4 or later.

Nonstop bridging is supported on EX Series switches with redundant Routing Engines in a Virtual Chassis or in a Virtual Chassis Fabric.

Nonstop bridging is supported on QFX Series switches in a Virtual Chassis or in a Virtual Chassis Fabric.

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



NOTE: All Routing Engines configured for nonstop bridging must be running the same Junos OS release.

Protocol Support

Nonstop bridging is supported for the following Layer 2 control protocols:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)
- VLAN Spanning Tree Protocol (VSTP)

Related Documentation

- [Nonstop Bridging Concepts on page 43](#)
- [Configuring Nonstop Bridging](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 49](#)

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 51](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 47](#)
- [Understanding Nonstop Active Routing on EX Series Switches](#)
- [Nonstop Active Routing Concepts on page 29](#)

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 46](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 51](#)
- [Example: Configuring Nonstop Active Routing](#)

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*
- [Nonstop Bridging Concepts on page 43](#)
- *Understanding In-Service Software Upgrade (ISSU)*

Resetting Local Statistics

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, accounting information, and trace options information are not replicated to the backup Routing Engine.

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. 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. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the **clear** command to clear statistics and protocol database information, see the [CLI Explorer](#).



NOTE: The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover.

Related Documentation

- [Understanding Graceful Routing Engine Switchover on page 13](#)
- [Configuring Graceful Routing Engine Switchover on page 23](#)

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 51](#)
- [Overview and Topology on page 51](#)
- [Configuration on page 52](#)
- [Verification on page 53](#)
- [Troubleshooting on page 53](#)

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 53](#)

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 53](#)

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 47](#).

**Related
Documentation**

- [Configuring Nonstop Active Routing on Switches on page 46](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 47](#)
- [*Understanding Nonstop Active Routing on EX Series Switches*](#)
- [Nonstop Active Routing Concepts on page 29](#)

PART 4

Virtual Router Redundancy Protocol

- [Using VRRP on page 57](#)

CHAPTER 4

Using VRRP

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 61](#)
- [Configuring the Startup Period for VRRP Operations on page 62](#)
- [Configuring the Advertisement Interval for the VRRP Master on page 62](#)
- [Configuring VRRP Preemption and Hold Time on page 63](#)
- [Configuring a Route to Be Tracked on page 64](#)
- [Configuring a Logical Interface to Be Tracked on page 65](#)
- [Configuring a Backup to Accept Packets Destined for the Virtual IP Address on page 67](#)
- [Configuring Passive ARP Learning for VRRP Backups on page 67](#)
- [Configuring the Silent Period on page 68](#)
- [Configuring Inheritance for a VRRP Group on page 68](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)
- [Troubleshooting VRRP on page 74](#)

Understanding VRRP

Juniper Networks switches support the Virtual Router Redundancy Protocol (VRRP) and VRRPv3 (for IPv6). This topic covers:

- [Overview of VRRP on page 57](#)
- [Sample VRRP Topology on page 58](#)

Overview of VRRP

Configuring end hosts on your network with static default routes minimizes configuration effort and complexity and reduces processing overhead on the end hosts. When hosts are configured with static routes, the failure of the default gateway normally results in a catastrophic event, isolating all hosts that are unable to detect available alternate paths to their gateway. Using Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for end hosts if the primary gateway fails.

VRRP (defined in RFC 3768) provides dynamic failover of IP addresses from one router to another in the event of failure. You can implement VRRP to provide a highly available default path to a gateway without needing to configure dynamic routing or router discovery protocols on end hosts.

Switches configured with VRRP share a virtual IP address, which is the address you configure as the default route on the hosts. At any time, one of the switches is the VRRP master, meaning that it owns the virtual IP address and is the active default gateway. The other devices are backups. The switches dynamically assign master and backup roles based on priorities that you configure (**1 through 255**). If the master fails, the backup switch with the highest priority becomes the master within a few seconds. This is done without any interaction with the hosts.

In VRRP operation, the master sends advertisements to the backup switches at regular intervals. The default interval is 1 second. If the backup switches do not receive an advertisement for a set period, the backup with the highest priority takes over as master within a few seconds and begins forwarding packets. This is done without any interaction with the hosts.




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

You can configure two QFabric systems to participate in a VRRP configuration as if they were two standalone switches. One benefit of this configuration is if you use VMware's vMotion, virtual machines can transition between hosts connected to the QFabric systems without updating their default gateway information. For example, a virtual machine running on a host connected to a QFabric system in data center A can transition to a host connected to a QFabric system in data center B without needing to resolve a new gateway IP address and MAC address.

Sample VRRP Topology

[Figure 7 on page 59](#) illustrates a basic VRRP topology. In this example, switches A and B are running VRRP and share the virtual IP address 10.1.1.1. The default gateway for each of the clients is 10.1.1.1.

Figure 7: Basic VRRP Topology



g041154

The following illustrates basic VRRP behavior using [Figure 7 on page 59](#) for reference:

1. When any of the servers wants to send traffic out of the LAN, it sends the traffic to the default gateway address of 10.1.1.1. This is a virtual IP address (VIP) owned by VRRP group 100. Because switch A is the master of the group, the VIP is associated with the “real” address 10.1.1.251 on switch A, and traffic from the servers is actually sent to this address. (Switch A is the master because it has been configured with a higher priority value.)
2. If there is a failure on switch A that prevents it from forwarding traffic to or from the servers—for example, if the interface connected to the LAN fails—switch B becomes the master and assumes ownership of the VIP. The servers continue to send traffic to the VIP, but because the VIP is now associated with the “real” address 10.1.1.252 on switch B (because of change of master), the traffic is sent to switch B instead of switch A.
3. If the problem that caused the failure on switch A is corrected, switch A becomes the master again and reasserts ownership of the VIP. In this case, the servers resume sending traffic to switch A.

Notice that no configuration changes are required on the servers for them to switch between sending traffic to switch A and switch B. When the VIP moves between 10.1.1.251 and 10.1.1.252, the change is detected by normal TCP-IP behavior and no configuration or intervention is required on the servers.

Related Documentation

- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Configuring Basic VRRP Support for QFX

To configure basic VRRP support, configure VRRP groups on interfaces by including the **vrrp-group** statement:

```
vrrp-group group-id {  
    priority number;  
    virtual-address [ addresses ];  
}
```

An interface can be a member of multiple VRRP groups.

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

For each interface, you must configure the following:

- Group identifier—Assign a value from 0 through 255. You must use the same identifier for each switch in the VRRP group.
- Priority—Assign a value from 1 through 255. The switch with the highest priority becomes the VRRP master. Assign different priorities to each switch in the VRRP group. If there are two or more switches with the same priority, the switch with the VRRP interface that has the highest IP address becomes the master.
- Virtual IP address—Normally, you configure only one address per group, but you can configure as many as eight addresses. Do not include a prefix length in a virtual IP address. The following considerations apply to configuring a virtual IP address:
 - You must configure the same address on all the switches in the VRRP group.
 - If you configure a virtual IP address to be the same as a physical interface address, the switch with that interface becomes the master for the group. You must configure the priority to be 255, and you must configure preemption by including the **preempt** statement.
 - If the virtual IP address is not the same as the physical interface address, you must ensure that the address does not appear anywhere else in the switch configuration. For example, verify that you do not use this address for another interface (including an aggregated Ethernet interface) or for a static ARP entry.



NOTE: If you enable MAC source address filtering on an interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement at the **[edit interfaces *interface-name*]** hierarchy. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 3768. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

- Related Documentation**
- [Understanding VRRP on page 57](#)
 - [Configuring the Startup Period for VRRP Operations on page 62](#)
 - [Configuring VRRP Authentication \(IPv4 Only\) on page 61](#)

Configuring VRRP Authentication (IPv4 Only)

VRRP (IPv4 only) protocol exchanges can be authenticated to guarantee that only trusted switches participate in a VRRP group. By default, VRRP authentication is disabled. You can configure one of the following authentication methods for a group, and each switch in the same group must use the same method:

- **Simple authentication**—Uses a text password included in the transmitted packet. The receiving switch uses an authentication key (password) to verify the packet.
- **Message Digest 5 (MD5) algorithm**—Adds an authentication header (AH) to the IP packet that encapsulates the VRRP packet. You create an authentication key that is used to create a hash of the packet, and the hash is stored in the AH. A receiving switch recalculates the hash on the incoming packet and compares the hashes. If they are identical, the packet is valid and is accepted. Otherwise the switch drops the incoming packet.

To enable authentication and specify an authentication method, include the **authentication-type** statement.

```
authentication-type authentication;
```

authentication can be **simple** or **md5**. The authentication type must be the same for all the switches in the VRRP group.

You can include this statement at the following hierarchy level:

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

If you include the **authentication-type** statement, you can configure a key (password) on each interface by including the **authentication-key** statement:

```
authentication-key key;
```

key (the password) is an ASCII string. For simple authentication, it can be from 1 through 8 characters long. For MD5 authentication, it can be from 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").



NOTE: The key must be the same for all switches in the VRRP group.

You can include this statement at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]`

**Related
Documentation**

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)

Configuring the Startup Period for VRRP Operations

Configure the startup-silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets while an interface is coming online. The period starts when the state of a VRRP interface is changed from down to up. During this period, Master Down Events are ignored.

To configure the startup period for VRRP operations, include the **startup-silent-period** statement at the `[edit protocols vrrp]` hierarchy level:

```
[edit protocols vrrp]  
  startup-silent-period seconds;
```

**Related
Documentation**

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Configuring the Advertisement Interval for the VRRP Master

By default, the master switch sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master switch is still operational. If the master switch fails or becomes unreachable, the backup switch with the highest priority value becomes the new master switch.

You can modify the advertisement interval in seconds or in milliseconds; the interval must be the same for all routing platforms in the VRRP group.

This topic contains the following sections:

- [Modifying the Advertisement Interval in Seconds on page 62](#)
- [Modifying the Advertisement Interval in Milliseconds on page 63](#)

Modifying the Advertisement Interval in Seconds

To modify the time, in seconds, between the sending of VRRP advertisement packets, include the **advertise-interval** statement:

```
  advertise-interval seconds;
```

The interval can be from 1 through 255 seconds.

You can include this statement at the following hierarchy level:

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

Modifying the Advertisement Interval in Milliseconds

To modify the time, in milliseconds, between the sending of VRRP advertisement packets, include the **fast-interval** statement:

```
fast-interval milliseconds;
```

The interval can be from 100 through 999 milliseconds.

You can include this statement at the following hierarchy level:

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



NOTE: Junos OS sets the advertisement interval to 0 in VRRP packets. When you configure VRRP with other vendors' equipment, the **fast-interval** statement works correctly only when the other equipment also has the advertisement interval set to 0 in the VRRP packet. Otherwise, Junos OS interprets other routers' settings as advertisement timer errors.

Related Documentation

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Configuring VRRP Preemption and Hold Time

- [Configuring VRRP Preemption on page 63](#)
- [Configuring the Preemption Hold Time on page 64](#)
- [Overriding the Hold Time on page 64](#)

Configuring VRRP Preemption

By default, a higher-priority VRRP backup switch preempts a lower-priority master switch. To explicitly enable this behavior, include the following statement:

```
preempt;
```

To prohibit a higher-priority VRRP backup switch from preempting a lower-priority master switch, include the following statement on the lower-priority switch:

```
no-preempt;
```

You can include these statements at the following hierarchy level:

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

Configuring the Preemption Hold Time

You can also configure a preemption hold time, which is the number of seconds a higher-priority backup router that has just started up waits before preempting the master router. You might want to configure a hold time so that routing protocols or other Junos OS components converge before preemption occurs.

The hold time is applied only on startup. By default, the hold-time value is 0 seconds, meaning that preemption can occur immediately after the backup router starts up.

To modify the preemption hold-time value, configure the following statement:

hold-time *seconds*;

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy level:

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

Overriding the Hold Time

You can use the **asymmetric-hold-time** statement to configure a VRRP master to fail over to the backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.

When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.

You can include this statement at the following hierarchy level:

- **[edit protocols vrrp]**

Related Documentation

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Configuring a Route to Be Tracked

A VRRP master can track a route and dynamically trigger a new master router election if the route becomes unreachable. To enable this behavior, you must configure a cost that will be subtracted from the priority of the master if the tracked route becomes

unreachable. The new priority must be less than the priority of one of the backups so that the backup becomes the new master.

To configure a route to be tracked, include the following statements:

```
track {
  priority-hold-time seconds;
  route prefix/prefix-length routing-instance default priority-cost priority;
}
```

You can include these statements at the following hierarchy level:

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

The **prefix** and **prefix-length** values specify the route to be tracked. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority of the master changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **priority-cost** option is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).

Related Documentation

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)
- [Configuring a Logical Interface to Be Tracked on page 65](#)

Configuring a Logical Interface to Be Tracked

VRRP can track whether a logical interface is up, down, or not present, and can change the priority of the switch based on the state of the interface, which might trigger a new master election. VRRP can also track the operational speed of a logical interface and update the priority of the switch when the speed crosses a configured threshold. For each VRRP group, you can track as many as 10 logical interfaces.

When interface tracking is enabled on a switch, you cannot assign the switch a priority of 255 to make it the master for the group.

To configure a logical interface to be tracked, include the following statements:

```
track {
  interface interface-name {
    bandwidth-threshold bits-per-second priority-cost priority;
    priority-cost priority;
  }
}
```

```

    priority-hold-time seconds;
}

```

You can include these statements at the following hierarchy level:

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

The interface specified is the interface to be tracked for the VRRP group. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **bandwidth-threshold** statement specifies a threshold for the tracked interface. If the bandwidth of the tracked interface drops below the threshold value, the system subtracts the bandwidth threshold **priority-cost** value from the VRRP priority for the switch. You can create as many as five **bandwidth-threshold** statements for each tracked interface.

The interface **priority-cost** statement is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).



WARNING: On a QFabric system, do not apply interface tracking to a multichassis link aggregation group (MC-LAG) that includes an interface belonging to a network Node group device and an interface belonging to a server Node group device. If you do apply interface tracking to an MC-LAG configured in this way, a priority update will not occur if the state of the MC-LAG interface changes.

If you configure tracking for more than one interface, Junos OS subtracts the sum of the priority costs for the tracked interfaces from the VRRP priority if all the tracked interfaces fail. However, if you configure the interface **priority-cost** statement and the bandwidth threshold **priority-cost** statement, they are not added together. The switch uses only one priority cost for a tracked interface, as indicated in [Table 7 on page 66](#):

Table 7: Interface State and Priority Cost Usage

Tracked Interface State	Priority Cost Usage
Down	priority cost <i>priority</i>
Not down; media speed below one or more bandwidth thresholds	Priority-cost of the lowest applicable bandwidth threshold

You must configure an interface priority cost only if you do not configure any bandwidth thresholds. If you do not configure an interface **priority-cost** value and the interface fails,

Junos OS subtracts the bandwidth threshold **priority-cost** value of the lowest bandwidth threshold from the priority of the switch.

**Related
Documentation**

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)
- [Configuring a Route to Be Tracked on page 64](#)

Configuring a Backup to Accept Packets Destined for the Virtual IP Address

By default, a switch configured to be a VRRP backup but acting as the master does not process packets sent to the virtual IP address—that is, packets in which the destination address is the virtual IP address. To configure a backup switch to process packets sent to the virtual IP address while it is acting as the master, include the **accept-data** statement on the backup:

accept-data;

You can include this statement at the following hierarchy level:

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

To explicitly prohibit the backup from accepting packets destined for the virtual IP address while acting as master, include the **no-accept-data** statement:

no-accept-data;

If you include the **accept-data** statement, configure the connected hosts so that they:

- Process gratuitous ARP requests.
- Do not use packets other than ARP replies to update their ARP cache.

This statement is disabled by default. If you enable it, your configuration does not comply with RFC 3768.

To restrict incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.

**Related
Documentation**

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Configuring Passive ARP Learning for VRRP Backups

By default, VRRP backup switches drop ARP requests for the MAC address of the VRRP IP. This means that backups do not learn the ARP mappings (IP address to MAC address mappings) for the hosts sending the requests. If it becomes the master, the configured backup must learn all the entries that were present in the ARP cache of the original

master. In environments with many directly attached hosts, the number of ARP entries to learn can be very large. This can cause a significant delay while the backup transitions to the master state, during which traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup to learn approximately the same contents as the ARP cache in the master, thus preventing the problem of needing to learn many ARP entries quickly. To enable passive ARP learning, include the **passive-learning** statement at the **[edit system arp]** hierarchy level:

```
[edit system arp]
passive-learning;
```

We recommend setting passive learning on both the backup and master VRRP switches. Doing so prevents the need to manually configure a master that fails and becomes a backup. While a switch operates as the master, the passive learning configuration has no impact. The configuration takes effect only when a switch operates as a backup.

**Related
Documentation**

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Configuring the Silent Period

When the state of a VRRP interface changes from down to up, a silent period begins. During this period, any master down events are ignored. Configure the silent period interval to avoid problems that can be caused if incoming VRRP advertisement packets are delayed or interrupted while an interface starts up.

To configure the silent period, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
startup-silent-period seconds;
```

**Related
Documentation**

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Configuring Inheritance for a VRRP Group

Junos OS enables you to configure VRRP groups on the various subnets of a VLAN to inherit the state and configuration of one of the groups, which is known as the *active VRRP group*. By configuring inheritance, you can prevent VRRP groups other than the active group from sending out VRRP advertisements. When the **vrrp-inherit-from** configuration statement is included in the configuration, only the active VRRP group from which the other VRRP groups are inheriting the state sends out VRRP advertisements; the groups inheriting the state do not send any VRRP advertisements, because the state is maintained only on the group from which the state is inherited.

If the **vrrp-inherit-from** statement is not configured, each of the VRRP master groups in the various subnets on the VLAN sends out separate VRRP advertisements and adds to the traffic on the VLAN.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]**:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
  vrrp-group group-id]
  vrrp-inherit-from vrrp-group;
```

When you configure a group to inherit a state from another group, note the following conditions:

- Both inheriting groups and active groups must be on the same physical interface and logical system. However, the groups need not necessarily be on the same VLAN or logical interface.
- Both inheriting groups and active groups must be on the same routing instances; however, this limitation does not apply for groups on the integrated routing and bridging (IRB) interfaces.

When you include the **vrrp-inherit-from** statement for a VRRP group, the VRRP group inherits the following parameters from the active group:

- **advertise-interval**
- **authentication-key**
- **authentication-type**
- **fast-interval**
- **preempt | no-preempt**
- **priority**
- **track interfaces**
- **track route**

However, you can configure the **accept-data | no-accept-data** statement for the group to specify whether the interface should accept packets destined for the virtual IP address.

Related Documentation

- [Understanding VRRP on page 57](#)
- [Configuring Basic VRRP Support for QFX on page 60](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Example: Configuring VRRP for Load Sharing

If you do not want to dedicate a switch to be a VRRP backup (and therefore leave it idle unless the master fails), you can create a load-sharing configuration in which each participating switch simultaneously acts as a master and a backup.

One reason to use a load-sharing (active-active) configuration is that you are more likely to actively monitor and maintain both switches and notice if a problem occurs on either of them. If you use a configuration in which one switch is only a backup (an active-backup configuration), you might be less likely to pay attention to the backup switch while it is idle. In the worst case, this could lead to the backup switch developing an undetected problem and not being able to perform adequately when a failover occurs.

- [Requirements on page 70](#)
- [Overview and Topology on page 70](#)
- [Configuring VRRP on Both Switches on page 71](#)
- [Verification on page 73](#)

Requirements

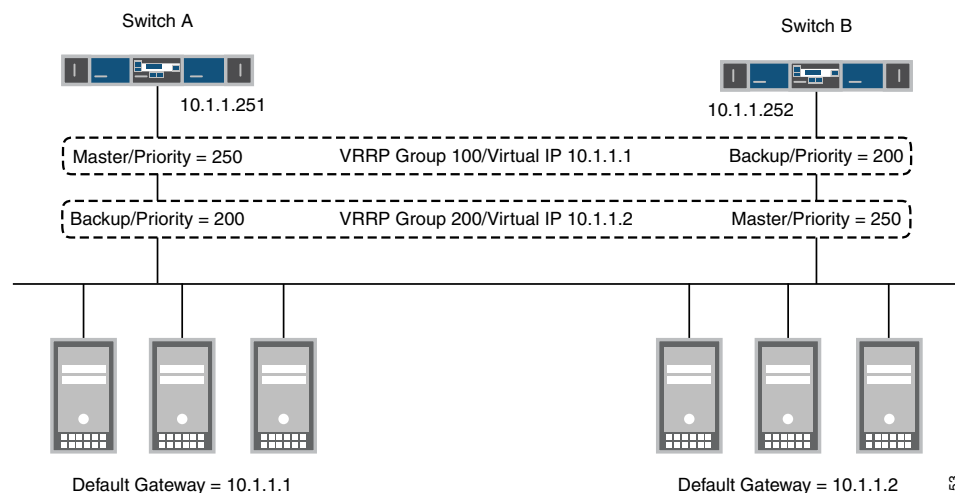
This example uses the following hardware and software components:

- Two switches
- Junos OS Release 11.3 or later
- Static routing or a dynamic routing protocol enabled on both switches.

Overview and Topology

This example uses two VRRP groups, each of which has its own virtual IP address. Devices on the LAN use one of these virtual IP addresses as their default gateway. If one of the switches fails, the other switch takes over for it. In the topology shown in [Figure 8 on page 70](#), for example, Switch A is the master for VRRP group 100. If Switch A fails, Switch B takes over and forwards traffic that the end devices send to the default gateway address 10.1.1.1.

Figure 8: VRRP Load-Sharing Configuration



g041153

This example shows a simple configuration to illustrate the basic steps for configuring two switches running VRRP to back each other up. [Table 8 on page 71](#) lists VRRP settings for each switch.

Table 8: Settings for VRRP Load-Sharing Example

Switch A	Switch B
VRRP Group 100: <ul style="list-style-type: none">Interface address: 10.1.1.251VIP: 10.1.1.1Priority: 250	VRRP Group 100: <ul style="list-style-type: none">Interface address: 10.1.1.252VIP: 10.1.1.1Priority: 200
VRRP Group 200: <ul style="list-style-type: none">Interface address: 10.1.1.251VIP: 10.1.1.2Priority: 200	VRRP Group 200: <ul style="list-style-type: none">Interface address: 10.1.1.252VIP: 10.1.1.2Priority: 250

In addition to configuring the two switches as shown, you must configure your end devices so that some of them use one of the virtual IP addresses as their default gateway and the remaining end devices use the other virtual IP address as their default gateway.

Note that if a failover occurs, the remaining switch might be unable to handle all of the traffic, depending on the demand.

Configuring VRRP on Both Switches

CLI Quick Configuration

```
Enter the following on Switch A:
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 priority 200
Enter the following on Switch B:
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 priority 250
```

Step-by-Step Procedure

- Configure the VRRP groups and priorities on Switch A:
- Create VRRP group 100 on Switch A and configure the virtual IP address for the group:
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
 - Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
100 priority 250
```

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
100 priority 200
```

Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch B:

1. Create VRRP group 100 on Switch B and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
100 priority 200
```

Switch A remains the master for group 100 because it has the highest priority for this group.

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
100 priority 250
```

Switch B becomes the master for group 200 because it has the highest priority for this group.

Results Display the results of the configuration on Switch A:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.251 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 250
          }
        }
      }
    }
  }
}
```



```
        vrrp-group 200 {
            virtual address 10.1.1.2
            priority 200
        }
    }
}
```

Display the results of the configuration on Switch B:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.252 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 200
          }
          vrrp-group 200 {
            virtual address 10.1.1.2
            priority 250
          }
        }
      }
    }
  }
}
```

Verification

- [Verifying that VRRP Is Working on Switch A on page 73](#)
- [Verifying that VRRP Is Working on Switch B on page 74](#)

Verifying that VRRP Is Working on Switch A

Purpose Verify that VRRP is active on Switch A and that the master and backup roles are correct.

Action Use the following command to verify that VRRP is active on Switch A and that the switch is master for group 100 and backup for group 200.

```
user@switch> show vrrp
```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	master	A .0327 1c1 10.1.1.251 vip 10.1.1.1	
xe-0/0/0.0	up	200	backup	A .0327 1c1 10.1.1.251 vip 10.1.1.2	

Meaning The `show vrrp` command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The `1c1` address is the physical address of the

interface and the **vip** address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 200 does not arrive before the timer expires, Switch A asserts itself as the master for this group.

Verifying that VRRP Is Working on Switch B

Purpose Verify that VRRP is active on Switch B and that the master and backup roles are correct.

Action Use the following command to verify that VRRP is active on Switch B and that the switch is backup for group 100 and master for group 200.

```
user@switch> show vrrp
```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	backup	A .0327 1c1	10.1.1.252 vip 10.1.1.1
xe-0/0/0.0	up	200	master	A .0327 1c1	10.1.1.252 vip 10.1.1.2

Meaning The **show vrrp** command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The **lcl** address is the physical address of the interface and the **vip** address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 100 does not arrive before the timer expires, Switch B asserts itself as the master for this group.

- Related Documentation**
- [Understanding VRRP on page 57](#)
 - [Configuring Basic VRRP Support for QFX on page 60](#)

Troubleshooting VRRP

Problem **Description:** If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).

Solution Configure a failover delay so that the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.

Related Documentation • [failover-delay on page 118](#)

PART 5

Nonstop Software Upgrade

- [Using NSSU on page 79](#)

CHAPTER 5

Using NSSU

- [Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade on page 79](#)

Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade

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

This topic covers:

- [Preparing the Switch for Software Installation on page 79](#)
- [Upgrading the Software Using NSSU on page 80](#)

Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

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

If the Virtual Chassis or mixed Virtual Chassis members are not running the same version of the software, use the **request system software add** command to upgrade the software on the inconsistent members.

- Ensure that graceful Routing Engine switchover (GRES) is enabled.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on each member to an external storage device with the **request system snapshot** command.

Upgrading the Software Using NSSU

This procedure describes how to upgrade the software running on all Virtual Chassis or mixed Virtual Chassis members using NSSU. When the upgrade completes, all members are running the new version of the software. Because a graceful Routing Engine switchover occurs during the upgrade, the original Virtual Chassis or mixed Virtual Chassis backup is the new master.



NOTE: Junos OS software images with enhanced automation are only supported on a non-mixed Virtual Chassis with QFX5100 switches. Also, performing an NSSU from a standard Junos OS software image to a Junos OS software image with enhanced automation or from a Junos OS software image with enhanced automation to a standard Junos OS software image is not supported.

To upgrade all members using NSSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Upgrading Software*.
2. Copy the software package or packages to the Virtual Chassis or mixed Virtual Chassis. We recommend that you copy the file to the **/var/tmp** directory on the master.
3. Log in to the Virtual Chassis or mixed Virtual Chassis using the console connection or the virtual management Ethernet (VME) interface. Using a console connection allows you to monitor the progress of the master switch reboot.
4. Start the NSSU:
 - On a Virtual Chassis or mixed Virtual Chassis, enter:

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

where **package-name.tgz** is, for example,
jinstall-qfx-3-13.2X50-D15.3-domestic-signed.tgz.

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

```
Chassis ISSU Check Done  
NSSU: Validating Image  
NSSU: Preparing Backup RE
```


Installing image on other FPC's along with the backup

Checking pending install on fpc1

Pushing bundle to fpc1

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc1

Checking pending install on fpc2

Pushing bundle to fpc2

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc2

Rebooting fpc1

NSSU: Backup RE Prepare Done

Waiting for Backup RE reboot

GRES operational

Initiating Chassis In-Service-Upgrade

Chassis NSSU Started

NSSU: Preparing Daemons

NSSU: Daemons Ready for NSSU

NSSU: Starting Upgrade for FRUs

NSSU: Preparing for Switchover

NSSU: Ready for Switchover

Checking In-Service-Upgrade status

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online (ISSU)	

Going to install image on master

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

relinquish mastership

NSSU: IDLE

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY

Shutdown NOW!

[pid 9336]

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

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

6. To ensure that the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partitions of all members:

```
user@switch> request system snapshot slice alternate all-members
```

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

Related Documentation

- *Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis*
- *Understanding Resilient Dual-Root Partitions on Switches*

PART 6

Configuration Statements and Operational Commands

- [Configuration Statements for Graceful Restart on page 85](#)
- [Configuration Statements for Graceful Switchover on page 101](#)
- [Configuration Statements for Nonstop Bridging and Routing on page 103](#)
- [Configuration Statements for VRRP on page 111](#)
- [Operational Commands for Graceful Restart on page 135](#)
- [Operational Commands for Graceful Switchover on page 165](#)
- [Operational Command for Nonstop Routing on page 175](#)
- [Operational Command for NSSU on page 179](#)
- [Operational Commands for VRRP on page 191](#)

CHAPTER 6

Configuration Statements for Graceful Restart

- [disable](#) on page 86
- [disable \(BGP Graceful Restart\)](#) on page 87
- [graceful-restart \(Enabling Globally\)](#) on page 88
- [graceful-restart \(Protocols BGP\)](#) on page 90
- [graceful-restart](#) on page 92
- [helper-disable \(OSPF\)](#) on page 94
- [no-strict-lsa-checking](#) on page 95
- [notify-duration](#) on page 96
- [redundancy \(Graceful Switchover\)](#) on page 97
- [restart-duration](#) on page 98
- [restart-time \(BGP Graceful Restart\)](#) on page 99
- [stale-routes-time](#) on page 100

disable

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (bgp isis ldp ospf ospf3 pim rip ripng rsvp) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (bgp ldp ospf ospf3 pim) graceful-restart],</p> <p>[edit protocols (bgp esis isis ospf ospf3 ldp pim rip ripng rsvp) graceful-restart],</p> <p>[edit protocols bgp group <i>group-name</i> graceful-restart],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>ip-address</i> graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (bgp ldp ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart],</p> <p>[edit routing-options graceful-restart]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</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	Disable graceful restart.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Graceful Restart</i> • Configuring Routing Protocols Graceful Restart on page 4 • <i>Configuring Graceful Restart for MPLS-Related Protocols</i> • <i>Configuring VPN Graceful Restart</i> • <i>Configuring Logical System Graceful Restart</i> • <i>Graceful Restart Configuration Statements</i> • <i>Configuring Graceful Restart for QFabric Systems</i>

disable (BGP Graceful Restart)


Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> graceful-restart], [edit protocols bgp graceful-restart], [edit protocols bgp group <i>group-name</i> graceful-restart], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Disable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition.



NOTE: When you disable graceful restart at one level in the configuration statement hierarchy, it is also disabled at lower levels in the same hierarchy. For example, if you disable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level, it is disabled for all the peers in the group. Therefore, if you want to enable graceful restart for some peers in a group and disable it for others, enable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level and disable graceful restart for each peer at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.


Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 5 • graceful-restart on page 90 • <i>restart-time</i> • <i>stale-routes-time</i>

graceful-restart (Enabling Globally)

Syntax	<pre> graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; notify-duration <i>seconds</i>; recovery-time <i>seconds</i>; restart-duration <i>seconds</i>; stale-routes-time <i>seconds</i>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p>
<div>  <p>NOTE:</p> <ul style="list-style-type: none"> For VPNs, the <code>graceful-restart</code> statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities. LDP sessions flap when <code>graceful-restart</code> configurations change. </div>	
Default	Graceful restart is disabled by default.
Options	The remaining statements are explained separately.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Enabling Graceful Restart Configuring Routing Protocols Graceful Restart on page 4

- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

graceful-restart (Protocols BGP)

Syntax	<pre> graceful-restart { disable; restart-time seconds; stale-routes-time seconds; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default. However, helper mode, the ability to assist a neighboring router attempting a graceful restart, is enabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <hr/> <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">  </div> <div> <p>NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> </div> <hr/> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 5 • Configuring Graceful Restart for QFabric Systems

- *Junos OS High Availability Library for Routing Devices*

graceful-restart

Syntax	<pre>graceful-restart { disable; helper-disable (standard restart-signaling both); no-strict-lsa-checking; notify-duration <i>seconds</i>; restart-duration <i>seconds</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the no-strict-lsa-checking statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode standard, restart-signaling, and both options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure graceful restart for OSPF.
Options	<p>disable—Disable graceful restart for OSPF.</p> <p>helper-disable (standard restart-signaling both)—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The standard, restart-signaling, and both options are only supported for OSPFv2. Specify standard to disable helper mode for standard graceful restart (based on RFC 3623). Specify restart-signaling to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify both to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p>Default: Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p>no-strict-lsa-checking—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols (ospf | ospf3)** command.

notify-duration seconds—Estimated time needed to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration seconds—Estimated time needed to reacquire a full OSPF neighbor from each area.


Range: 1 through 3600 seconds

Default: 180 seconds

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

Related Documentation	• <i>Example: Configuring Graceful Restart for OSPF</i>
	• <i>Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart</i>
	• <i>Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart</i>
	• <i>Example: Disabling Strict LSA Checking for OSPF Graceful Restart</i>

helper-disable (OSPF)

Syntax	<code>helper-disable < both restart-signaling standard >;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf graceful-restart], [edit protocols ospf graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ospf graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Options both , restart-signaling , and standard introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Disable helper mode for graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart. The last committed statement takes precedence over the previously configured statement.
Default	Helper mode is enabled by default for OSPF.
Options	both —(Optional) Disable helper mode for both standard and restart signaling-based graceful restart. restart-signaling —(Optional) Disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813).
<div>  <p>NOTE: Restart signaling-based helper mode is not supported for OSPFv3 configurations.</p> </div>	
	standard —(Optional) Disable helper mode for standard graceful restart (based on RFC 3623).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring Routing Protocols Graceful Restart on page 4 Configuring Graceful Restart for MPLS-Related Protocols

no-strict-lsa-checking

Syntax	no-strict-lsa-checking;
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router or switch.
Default	By default, LSA checking is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for OSPF and OSPFv3 on page 7 • <i>Configuring Graceful Restart for QFabric Systems</i> • <i>maximum-neighbor-recovery-time</i> • <i>recovery-time</i>

notify-duration

Syntax	<code>notify-duration seconds;</code>
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart], [edit routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the length of time the router or switch notifies helper OSPF routers that it has completed graceful restart.
Options	seconds —Length of time in the router notifies helper OSPF routers that it has completed graceful restart. Range: 1 through 3600 Default: 30
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for OSPF and OSPFv3 on page 7• Configuring Graceful Restart for QFabric Systems• restart-duration on page 98

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 101 • Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 25 • Configuring Graceful Routing Engine Switchover on page 23 • Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure) • High Availability Features for EX Series Switches Overview

restart-duration

Syntax	<code>restart-duration seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (isis ospf ospf3 pim) graceful-restart],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</code> <code>[edit protocols (esis isis ospf ospf3 pim) graceful-restart],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</code> <code>[edit routing-options graceful-restart]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure the grace period for graceful restart globally.</p> <p>Additionally, you can individually configure the duration of the graceful restart period for the End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and OSPFv3 protocols and for Protocol Independent Multicast (PIM) sparse mode.</p>
Options	<p>seconds—Time for the graceful restart period.</p> <p>Range:</p> <p>The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none">• [edit routing-options graceful-restart] (global setting)—120 through 900• ES-IS—30 through 300• IS-IS—30 through 300• OSPF/OSPFv3—1 through 3600• PIM—30 through 300 <p>Default:</p> <p>The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none">• [edit routing-options graceful-restart] (global setting)—300• ES-IS—180• IS-IS—210• OSPF/OSPFv3—180• PIM—60

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

Related Documentation

- *Enabling Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Graceful Restart for VPNs*
- *Configuring Logical System Graceful Restart*

restart-time (BGP Graceful Restart)

Syntax restart-time *seconds*;

Hierarchy Level [edit protocols (bgp | rip | ripng) [graceful-restart](#)],
[edit logical-systems *logical-system-name* protocols (bgp | rip | ripng) [graceful-restart](#) ([Enabling Globally](#))],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp [graceful-restart](#)],
[edit routing-instances *routing-instance-name* protocols bgp [graceful-restart](#)]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.

Options *seconds*—Length of time for the graceful restart period.
Range: 1 through 600 seconds
Default: Varies by protocol:

- BGP—120 seconds
- RIP and RIPng—60 seconds

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

Related Documentation

- [Configuring Graceful Restart Options for BGP on page 5](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 9](#)
- *Configuring Graceful Restart for QFabric Systems*
- [stale-routes-time on page 100](#)

stale-routes-time

Syntax	<code>stale-routes-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-routing-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Specify the maximum time that stale routes are kept during a restart. The stale-routes-time statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.
Options	seconds —Time the router device waits to receive messages from restarting neighbors before declaring them down. Range: 1 through 600 seconds Default: 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP on page 5• Configuring Graceful Restart for QFabric Systems• restart-time (BGP Graceful Restart) on page 99

CHAPTER 7

Configuration Statements for Graceful Switchover

- [graceful-switchover](#) on page 101
- [redundancy \(Graceful Switchover\)](#) on page 102

[graceful-switchover](#)

Syntax	<code>graceful-switchover;</code>
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 51• Configuring Graceful Routing Engine Switchover on page 23• Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 25• Configuring Nonstop Active Routing on Switches on page 46• Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)

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 101• Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 25• Configuring Graceful Routing Engine Switchover on page 23• Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)• High Availability Features for EX Series Switches Overview

CHAPTER 8

Configuration Statements for Nonstop Bridging and Routing

- [nonstop-bridging on page 103](#)
- [nonstop-routing on page 104](#)
- [synchronize on page 105](#)
- [traceoptions on page 107](#)

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 49

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

- *Configuring Nonstop Active Routing*

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*

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 • *Example: Tracing Global Routing Protocol Operations*
Documentation

CHAPTER 9


Configuration Statements for VRRP

- [accept-data](#) on page 112
- [advertise-interval](#) on page 113
- [asymmetric-hold-time](#) on page 114
- [authentication-key](#) on page 115
- [authentication-type](#) on page 116
- [bandwidth-threshold](#) on page 117
- [failover-delay](#) on page 118
- [fast-interval](#) on page 119
- [hold-time \(VRRP\)](#) on page 120
- [interface](#) on page 121
- [preempt \(VRRP\)](#) on page 122
- [priority \(Protocols VRRP\)](#) on page 123
- [priority-cost \(VRRP\)](#) on page 124
- [priority-hold-time](#) on page 125
- [route \(Interfaces\)](#) on page 126
- [startup-silent-period](#) on page 127
- [traceoptions](#) on page 128
- [track \(VRRP\)](#) on page 130
- [virtual-address](#) on page 131
- [vrrp-group](#) on page 132

accept-data

Syntax	(accept-data no-accept-data);
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[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>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not an interface accepts packets destined for the virtual IP address:</p> <ul style="list-style-type: none"> • accept-data—Enable the interface to accept packets destined for the virtual IP address. • no-accept-data—Prevent the interface from accepting packets destined for the virtual IP address.
Default	<p>If the accept-data statement is not configured, the master router responds to ARP requests only.</p> <p>The accept-data statement has the following restrictions and limitations:</p> <ul style="list-style-type: none"> • If the master router owns the virtual IP address or if the priority of the master router is set to 255, the accept-data statement becomes inapplicable. • If the master router owns the virtual IP address, the master router responds to Internet Control Message Protocol (ICMP) message requests. • If you want to restrict the incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets. • If you include the accept-data statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, <i>Virtual Router Redundancy Protocol (VRRP)</i>).
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"> • <i>Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group</i>


advertise-interval

Syntax	<code>advertise-interval seconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
<div>  <p>NOTE: When VRRPv3 is enabled, the <code>advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</p> </div>	
Options	<p>seconds—Interval between advertisement packets.</p> <p>Range: 1 through 255 seconds</p> <p>Default: 1 second</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"> • <i>Configuring the Advertisement Interval for the VRRP Master Router</i> • fast-interval on page 119 • <i>inet6-advertise-interval</i> • <i>version-3</i>


asymmetric-hold-time

Syntax	asymmetric-hold-time;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	<p>Configure a VRRP master to fail over to a backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.</p> <p>When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.</p>
Default	asymmetric-hold-time 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">• Configuring VRRP Preemption and Hold Time on page 63

authentication-key

Syntax	<code>authentication-key key;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS 12.3X48-D10 for the SRX Series.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement. All routers in the VRRP group must use the same authentication scheme and password.
<div>  NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups. </div>	
Options	key —Authentication password. For simple authentication, it can be 1 through 8 characters long. For Message Digest 5 (MD5) authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").
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 Authentication (IPv4 Only) • Configuring VRRP Authentication (IPv4 Only) on page 61 • authentication-type on page 116 • version-3

authentication-type

Syntax	<code>authentication-type <i>authentication</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the authentication-key statement.</p> <p>All routers in the VRRP group must use the same authentication scheme and password.</p>
<div>  <p>NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups.</p> </div>	
Options	<p><i>authentication</i>—Authentication scheme:</p> <ul style="list-style-type: none"> simple—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure. md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme. <p>Default: none (no authentication is performed).</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 Authentication (IPv4 Only) • Configuring VRRP Authentication (IPv4 Only) on page 61 • authentication-key on page 115 • version-3


bandwidth-threshold

Syntax	<code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</p> <p>[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> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
Options	<p><i>bits-per-second</i>—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.</p> <p>Range: 1 through 10000000000000 bits per second</p> <p><i>priority-cost priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, and forces a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</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 a Logical Interface to Be Tracked for a VRRP Group • Configuring a Logical Interface to Be Tracked on page 65

failover-delay

Syntax	<code>failover-delay <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	<p>If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).</p> <p>If you configure a failover delay, the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.</p>
Options	<i>milliseconds</i> —Specify the failover delay time, in milliseconds. Range: 50 through 2000
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">• Troubleshooting VRRP on page 74• show vrrp on page 192

fast-interval

Syntax	<code>fast-interval milliseconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[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>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS 12.3X48-D10 for the SRX Series.</p>
Description	<p>Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
Options	<p><i>milliseconds</i>—Interval between advertisement packets.</p> <p>Range: 10 through 40,950 milliseconds (range extended from 100–999 to 10–40,950 in Junos OS Release 12.2).</p>
<div>  <p>NOTE: When configuring VRRP for IPv4, if you have chosen not to enable VRRPv3, you cannot set a value less than 100 for <i>fast-interval</i>. Commit check fails if a value less than 100 is configured.</p> </div>	
Default: 1 second	
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 the Advertisement Interval for the VRRP Master Router • Configuring the Advertisement Interval for the VRRP Master on page 62 • advertise-interval on page 113 • advertise-interval on page 113 • <i>inet6-advertise-interval</i> • <i>version-3</i>


hold-time (VRRP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.
Default	VRRP preemption is not timed.
Options	<i>seconds</i> —Hold-time period. Range: 0 through 3600 seconds Default: 0 seconds (VRRP preemption is not timed.)
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 a Backup Router to Preempt the VRRP Master Router• Configuring VRRP Preemption and Hold Time on page 63

interface

Syntax	<pre>interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> <i>priority-cost</i> <i>priority</i>; priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</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> track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>bandwidth-threshold statement added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p><i>interface-name</i>—Interface to be tracked for this VRRP group.</p> <p>Range: 1 through 10 interfaces</p> <p>The remaining statements are described 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"> • <i>Configuring a Logical Interface to Be Tracked for a VRRP Group</i> • Configuring a Logical Interface to Be Tracked on page 65 • <i>Junos OS Services Interfaces Library for Routing Devices</i>

preempt (VRRP)

Syntax	(preempt no-preempt) { hold-time <i>seconds</i> ; }
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>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> 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 before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a master router: <ul style="list-style-type: none"> • preempt—Allow the master router to be preempted. <p>.....</p> <div>  NOTE: By default, a higher-priority backup router can preempt a lower-priority master router. </div> <p>.....</p> <ul style="list-style-type: none"> • no-preempt—Prohibit the preemption of the master router. When no-preempt is configured, the backup router cannot preempt the master router even if the backup router has a higher priority. <p>The remaining statement is explained separately.</p>
Default	By default the preempt statement is enabled, and a higher-priority backup router preempts a lower-priority master router even if the preempt statement is not explicitly configured.
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 a Backup Router to Preempt the VRRP Master Router</i> • Configuring VRRP Preemption and Hold Time on page 63


priority (Protocols VRRP)

Syntax	<code>priority <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[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>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	<p>priority—Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p>Range: 1 through 255</p> <p>Default: 100 (for backup routers)</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 Basic VRRP Support • Configuring Basic VRRP Support for QFX on page 60

priority-cost (VRRP)

Syntax	<code>priority-cost priority;</code>
Hierarchy Level	<code>[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series. Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	priority —The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group. Range: 1 through 254
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 a Logical Interface to Be Tracked for a VRRP Group• Configuring a Logical Interface to Be Tracked on page 65

priority-hold-time

Syntax	<code>priority-hold-time seconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.</p>
<div>  <p>NOTE: When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP master will pre-empt according to the configuration of the priority-hold time.</p> </div>	
Options	<p>seconds—Minimum length of time that must elapse between dynamic priority changes.</p> <p>Range: 0through 3600 seconds</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 a Logical Interface to Be Tracked for a VRRP Group • Configuring a Logical Interface to Be Tracked on page 65

route (Interfaces)

Syntax	<code>route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS 11.3 for QFX Series. Statement introduced in Junos OS 12.1 for EX Series switches.
Description	Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p><i>prefix</i>—Route to be tracked for this VRRP group.</p> <p><i>priority-cost priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p> <p><i>routing-instance instance-name</i>—Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for <i>instance-name</i> must be default.</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 a Route to Be Tracked for a VRRP Group• Configuring a Route to Be Tracked on page 64

startup-silent-period

Syntax	<code>startup-silent-period <i>seconds</i>;</code>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Instruct the system to ignore the Master Down Event when an interface transitions from the down state to the up state. This statement is used to avoid incorrect error alarms caused by the delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase.
Options	<i>seconds</i> —Number of seconds for the startup period. Default: 4 seconds Range: 1 through 2000 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Startup Period for VRRP Operations</i>• Configuring the Startup Period for VRRP Operations on page 62

traceoptions

Syntax traceoptions {
 file <filename> <files number> <match *regular-expression*> <microsecond-stamp>
 <size size> <world-readable | no-world-readable>;
 flag *flag*;
 no-remote-trace;
 }

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

Description Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.

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

By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory **/var/log**.



NOTE: The traceoptions statement is not supported on a QFabric system.

Default If you do not include this statement, no VRRP-specific tracing operations are performed.

Options **filename filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, VRRP tracing output is placed in the file **vrrpd**.

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. When the maximum number is reached, the oldest trace file is overwritten.

Range: 0 through 4,294,967,296 files

Default: 3 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the VRRP-specific tracing options:

- **all**—All VRRP tracing operations
- **database**—Database changes
- **general**—General events

- **interfaces**—Interface changes
- **normal**—Normal events
- **packets**—Packets sent and received
- **state**—State transitions
- **timer**—Timer events

match *regex*—(Optional) Refine the output to include only those lines that match the given regular expression.

microsecond-stamp—(Optional) Provide a timestamp with microsecond granularity.

no-world-readable—Restrict users from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. 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.

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 routing platform

Default: 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

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>Tracing VRRP Operations</i>

track (VRRP)

Syntax	<pre>track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<pre>[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>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>priority-hold-time statement added in Junos OS Release 8.1.</p> <p>route statement added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.
Options	The remaining statements are described separately.
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 a Logical Interface to Be Tracked for a VRRP Group• Configuring a Route to Be Tracked for a VRRP Group• Configuring a Logical Interface to Be Tracked on page 65• Configuring a Route to Be Tracked on page 64

virtual-address

Syntax	<code>virtual-address [<i>addresses</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.
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"> • <i>Configuring Basic VRRP Support</i> • Configuring Basic VRRP Support for QFX on page 60

vrrp-group

Syntax	<pre> vrrp-group <i>group-id</i> { (accept-data no-accept-data); advertise-interval <i>seconds</i>; advertisements-threshold <i>number</i>; authentication-key <i>key</i>; authentication-type <i>authentication</i>; fast-interval <i>milliseconds</i>; (preempt no-preempt) { hold-time <i>seconds</i>; } priority <i>number</i>; track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; } virtual-address [<i>addresses</i>]; vrrp-inherit-from <i>vrrp-group</i>; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group. As of Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the nonstop-routing statement at the [edit routing-options] or [edit logical system <i>logical-system-name</i> routing-options hierarchy level.</p>
Options	<p><i>group-id</i>—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:53:00 through 00:00:5e:00:53:ff are reserved for VRRP, as defined in RFC 2338. 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	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Configuring Basic VRRP Support*
 - *Configuring VRRP*
 - [Configuring Basic VRRP Support for QFX on page 60](#)
 - [Example: Configuring VRRP for Load Sharing on page 69](#)
 - *vrrp-inet6-group*
 - [nonstop-routing on page 104](#)

CHAPTER 10

Operational Commands for Graceful Restart

- [Verifying Graceful Restart Operation on page 135](#)
- [show bgp neighbor](#)
- [show log](#)
- [show \(ospf | ospf3\) overview](#)

Verifying Graceful Restart Operation

This topic contains the following sections:

- [Graceful Restart Operational Mode Commands on page 135](#)
- [Verifying BGP Graceful Restart on page 136](#)
- [Verifying IS-IS and OSPF Graceful Restart on page 136](#)
- [Verifying CCC and TCC Graceful Restart on page 137](#)

Graceful Restart Operational Mode Commands

To verify proper operation of graceful restart, use the following commands:

- **show bgp neighbor** (for BGP graceful restart)
- **show log** (for IS-IS and OSPF/OSPFv3 graceful restart)
- **show (ospf | ospfv3) overview** (for OSPF/OSPFv3 graceful restart)
- **show rsvp neighbor detail** (for RSVP graceful restart—helper router)
- **show rsvp version** (for RSVP graceful restart—restarting router)
- **show ldp session detail** (for LDP graceful restart)
- **show connections** (for CCC and TCC graceful restart)
- **show route instance detail** (for Layer 3 VPN graceful restart and for any protocols using graceful restart in a routing instance)
- **show route protocol l2vpn** (for Layer 2 VPN graceful restart)

For more information about these commands and a description of their output fields, see the [CLI Explorer](#).

Verifying BGP Graceful Restart

To view graceful restart information for BGP sessions, use the **show bgp neighbor** command:

```
user@PE1> show bgp neighbor 192.0.2.10
Peer: 192.0.2.10+179 AS 64496 Local: 192.0.2.5+1106 AS 64496
  Type: Internal    State: Established    Flags: <>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ static ]
  Options:<Preference LocalAddress HoldTime GracefulRestart Damping PeerAS Refresh>

  Local Address: 192.168.5.1 holdtime: 90 Preference: 170
  IPsec SA Name: hope
  Number of flaps: 0
  Peer ID: 192.0.2.10    Local ID: 192.0.2.5    Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 180
  Stale routes from peer are kept for: 180
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
  Last traffic (seconds): Received 19    Sent 19    Checked 19
  Input messages: Total 2    Updates 1    Refreshes 0    Octets 42
  Output messages: Total 3    Updates 0    Refreshes 0    Octets 116
  Output Queue[0]: 0
```

Verifying IS-IS and OSPF Graceful Restart

To view graceful restart information for IS-IS and OSPF, configure traceoptions (see [“Tracking Graceful Restart Events” on page 10](#)).

Here is the output of a traceoptions log from an OSPF restarting router:

```
Oct 8 05:20:12 Restart mode - sending grace lsas
Oct 8 05:20:12 Restart mode - estimated restart duration timer triggered
Oct 8 05:20:13 Restart mode - Sending more grace lsas
```

Here is the output of a traceoptions log from an OSPF helper router:

```
Oct 8 05:20:14 Helper mode for neighbor 192.0.2.5
Oct 8 05:20:14 Received multiple grace lsa from 192.0.2.5
```


Verifying CCC and TCC Graceful Restart

To view graceful restart information for CCC and TCC connections, use the **show connections** command. The following example assumes four remote interface CCC connections between CE1 and CE2:

```
user@PE1> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP
```

CCC Graceful restart : Restarting

Connection/Circuit	Type	St	Time last up	# Up trans
CE1-CE2-0	rmt-if	Restart	-----	0
fe-1/1/0.0	intf	Up		
PE1-PE2-0	tlsp	Up		
PE2-PE1-0	rlsp	Up		
CE1-CE2-1	rmt-if	Restart	-----	0
fe-1/1/0.1	intf	Up		
PE1-PE2-1	tlsp	Up		
PE2-PE1-1	rlsp	Up		
CE1-CE2-2	rmt-if	Restart	-----	0
fe-1/1/0.2	intf	Up		
PE1-PE2-2	tlsp	Up		
PE2-PE1-2	rlsp	Up		
CE1-CE2-3	rmt-if	Restart	-----	0
fe-1/1/0.3	intf	Up		
PE1-PE2-3	tlsp	Up		
PE2-PE1-3	rlsp	Up		

- Related Documentation**
- [Graceful Restart Concepts on page 3](#)
 - *Configuring Graceful Restart for QFabric Systems*

show bgp neighbor

List of Syntax	Syntax on page 138 Syntax (EX Series Switch, QFX Series, and OCX Series) on page 138
Syntax	<pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (detail <i>neighbor-address</i>)</pre>
Syntax (EX Series Switch, QFX Series, and OCX Series)	<pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <neighbor-address> <orf (<i>neighbor-address</i> detail)</pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 14.1x53-D20 for the OCX Series. orf option introduced in Junos OS Release 9.2. exact-instance option introduced in Junos OS Release 11.4.
Description	Display information about BGP peers.
Options	<p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p>
Additional Information	For information about the local-address , nlri , hold-time , and preference statements, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .
Required Privilege Level	view

Related Documentation • [clear bgp neighbor](#)

List of Sample Output [show bgp neighbor on page 146](#)
[show bgp neighbor \(CLNS\) on page 147](#)
[show bgp neighbor \(Layer 2 VPN\) on page 148](#)
[show bgp neighbor \(Layer 3 VPN\) \(Not supported on the OCX Series.\) on page 150](#)
[show bgp neighbor neighbor-address on page 151](#)
[show bgp neighbor neighbor-address on page 151](#)
[show bgp neighbor neighbor-address \(BGP Graceful Restart Enabled\) on page 152](#)
[show bgp neighbor neighbor-address \(BGP Long-Lived Graceful Restart\) on page 153](#)
[show bgp neighbor orf neighbor-address detail on page 153](#)

Output Fields [Table 9 on page 139](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 9: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .
State	Current state of the BGP session: <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer. • route reflector client—The BGP session is established with a route reflector client.

Table 9: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • RSync—This peer in the backup Routing Engine is synchronized with the BGP peer in the master Routing Engine for nonstop active routing. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.

Table 9: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.
Options	<p>Configured BGP options:</p> <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AdvertiseBGPStatic—Configured BGP static routes are advertised. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • LLGR—BGP long-lived graceful restart capability is configured. • LLGRHelperDisabled—BGP long-lived graceful restart is completely disabled for a neighbor. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group.
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Peer does not support LLGR Restarter or Receiver functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode completely.

Table 9: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Peer does not support LLGR Restarter functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode for any family.
Authentication key change	(appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .
Address families configured	Names of configured address families for the VPN.
BGP-Static Advertisement Policy	Name of the bgp static policy that is configured on the peer.
Local Address	Address of the local routing device.
Remove-private options	Options associated with the remove-private statement.
Holdtime	Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the preference statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the out-delay parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.

Table 9: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGp peering is established.
NLRI and times for LLGR configured on peer	<p>Names of address families and stale time for BGP long-lived graceful restart configured on the BGP peer or neighbor.</p> <p>Times are displayed using the routing protocol daemon (rpd) %#OT format:</p> <p><weeks>w<days>d <hours>:<minutes>:<seconds></p> <p>Zero leading elements are omitted, for example, a value less than one week do not include the weeks.</p>
NLRI and times that peer supports LLGR Restarter for	<p>Names of address families and stale time that the BGP peer supports for restarter mode for BGP long-lived graceful restart.</p> <p>Times are displayed using the routing protocol daemon (rpd) %#OT format:</p> <p><weeks>w<days>d <hours>:<minutes>:<seconds></p> <p>Zero leading elements are omitted, for example, a value less than one week do not include the weeks.</p>
NLRI that peer saved LLGR forwarding for	Name of the address family for which the BGP peer saved BGP long-lived graceful restart forwarding.
Graceful Restart Details	Amount of time that is remaining until LLGR expires and the time remaining on the GR stale timer, along with RIB details, are displayed while LLGR receiver mode is active (a peer that negotiated LLGR has disconnected and not yet reconnected)
NLRI we are holding stale routes for	Names of address families (NLRIs) for which that stale routes are held or preserved when BGP graceful restart receiver mode is active for a neighbor.
Time until end-of-rib is assumed for stale routes	<p>Amount of time remaining on the stale timer until which end-of-RIB (EoR) markers are assumed when BGP graceful restart receiver mode is active for a neighbor.</p> <p>Time is displayed in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). Note that the stale timer display ('Time until end-of-rib is assumed') is also present when a session is active, but the neighbor as not yet sent all of the end-of-rib indications.</p>
Time until stale routes are deleted or become long-lived stale	Amount of time up to which stale routes are deleted or become long-lived stale routes when BGP graceful restart receiver mode is active for a neighbor.

Table 9: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as 1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.

Table 9: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast .
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast .
Table inet.number	Information about the routing table: <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	Information about dropped path attributes: <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	Information about ignored path attributes: <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.

Table 9: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Output queue	<p>Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.</p> <p>It also specifies the routing table name and the NLRI that the table was advertised through, in the format (<i>routing table name, NLRI</i>).</p> <p>NOTE: The output queue of routing tables that are not advertised, will only show up at extensive output level.</p>
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates rcv	<p>(orf option only) Number of outbound-route filters received for each configured address family.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>
Immediate	<p>(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redistrib_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Options: <AdvertiseBGPStatic>

```

```

Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
Number of flaps: 0
Peer ID: 10.255.7.250    Local ID: 10.255.7.248    Active Holdtime: 90
Keepalive Interval: 30    Group index: 0    Peer index: 0
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 10)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages: Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0 (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214    Local ID: 10.255.167.205    Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0    Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1    Local ID: 10.245.245.3    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
  Table bgp.isovpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          3

```

```

Received prefixes:          3
Suppressed due to damping:  0
Advertised prefixes:       3
Table aaaa.iso.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes:           3
Received prefixes:         3
Suppressed due to damping:  0
Last traffic (seconds): Received 6    Sent 5    Checked 5
Input messages: Total 1736    Updates 4    Refreshes 0    Octets 33385
Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305
Output Queue[0]: 0 (bgp.isovpn.0, iso-vpn-unicast)
Output Queue[1]: 0 (aaaa.iso.0, iso-vpn-unicast)

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65536 Local: 10.69.103.1      AS 65539
Type: External      State: Active      Flags: <ImportEval>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ BGP-INET-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-unicast
Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.69.104.2      AS 65536 Local: 10.69.104.1      AS 65539
Type: External      State: Active      Flags: <ImportEval>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ BGP-L-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-labeled-unicast
Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
Type: Internal      State: Established      Flags: <ImportEval>
Last State: OpenConfirm      Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast l2vpn
Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn

```

```

NLRI of received end-of-rib markers: inet-vpn-unicast 12vpn
Table bgp.13vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.12vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table L2VPN.12vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync

```

```

Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (bgp.l2vpn.0, inet-vpn-unicast)
Output Queue[2]: 0 (BGP-INET.inet.0, inet-vpn-unicast)
Output Queue[3]: 0 (BGP-L.inet.0, inet-vpn-unicast)
Output Queue[4]: 0 (LDP.inet.0, inet-vpn-unicast)
Output Queue[5]: 0 (OSPF.inet.0, inet-vpn-unicast)
Output Queue[6]: 0 (RIP.inet.0, inet-vpn-unicast)
Output Queue[7]: 0 (STATIC.inet.0, inet-vpn-unicast)
Output Queue[8]: 0 (L2VPN.l2vpn.0, inet-vpn-unicast)

```

show bgp neighbor (Layer 3 VPN) (Not supported on the OCX Series.)

```

user@host> show bgp neighbor
Peer: 192.0.2.0+179    AS 10045 Local: 192.0.2.1+1214    AS 10045
Type: Internal    State: Established    Flags: <ImportEval>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Export: [ match-all ] Import: [ match-all ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 192.0.2.1 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
                                size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110    Local ID: 192.168.1.111    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2

```

```

    Suppressed due to damping: 0
    Last traffic (seconds): Received 15    Sent 20    Checked 20
    Input messages: Total 40    Updates 2    Refreshes 0    Octets 856
    Output messages: Total 44    Updates 2    Refreshes 0    Octets 1066
    Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
    Output Queue[1]: 0 (vpn-green.inet.0, inet-vpn-unicast)
    Trace options: detail packets
    Trace file: /var/log/bgpgr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
  Keepalive Interval: 30
  BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3    Sent 3    Checked 3
  Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
  Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
  Output Queue[0]: 0 (inet.0, inet-unicast)
  Output Queue[1]: 0 (inet6.0, inet6-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpgr size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External    State: Established    Flags: <Sync>

```

```

Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6      Local ID: 10.255.245.5      Active Holdtime: 60000
Keepalive Interval: 20000   Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet.2, inet-multicast)
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

show bgp neighbor neighbor-address (BGP Graceful Restart Enabled)

```
user@router> show bgp neighbor 10.255.255.16
```

```

Peer: 10.255.255.16 AS 100      Local: 10.255.255.12 AS 100
Type: Internal   State: Active   Flags: <>
Last State: Idle   Last Event: Start
Last Error: None
Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
Options: <LLGR>
Address families configured: 12vpn
Local Address: 10.255.255.12 Holdtime: 90 Preference: 170
NLRI 12vpn:
Number of flaps: 6
Last flap event: Restart
NLRI we are holding stale routes for: inet-vpn-unicast
Time until stale routes are deleted or become long-lived stale: 00:01:57
Time until end-of-rib is assumed for stale routes: 00:04:43
Table bgp.13vpn.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete

```



```

Send state: not advertising
Active prefixes:          0
Received prefixes:        7
Accepted prefixes:        7
Suppressed due to damping: 0
Table foo.inet.0 Bit: 30000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not in sync
Active prefixes:          0
Received prefixes:        7
Accepted prefixes:        7
Suppressed due to damping: 0

```

show bgp neighbor neighbor-address (BGP Long-Lived Graceful Restart)

```

user@router> show bgp neighbor 10.4.12.11

Peer: 10.4.12.11 AS 100          Local: 10.6.128.225 AS 100
Type: Internal      State: Active      Flags: <>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ foo ]
Options: <Preference LocalAddress Refresh GracefulRestart>
Options: <LLGR>
Local Address: 10.6.128.225 Holdtime: 90 Preference: 170
Number of flaps: 3
Last flap event: Restart
Error: 'Cease' Sent: 0 Recv: 1
Time until long-lived stale routes deleted: inet-vpn-unicast 10:00:22
route-target 10:00:22
Table bgp.l3vpn.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes:          0
Received prefixes:        7
Accepted prefixes:        7
Suppressed due to damping: 0
Table foo.inet.0 Bit: 30000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not in sync
Active prefixes:          0
Received prefixes:        7
Accepted prefixes:        7
Suppressed due to damping: 0

```

show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
Filter updates rcv:          1 Immediate:          1
Filter: prefix-based receive
Received filter entries:
seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast

```

```
Filter updates rcv:          0 Immediate:          1
Filter: prefix-based receive
Received filter entries:
*:*
```

show log

List of Syntax	Syntax on page 155 Syntax (QFX Series and OCX Series) on page 155 Syntax (TX Matrix Router) on page 155
Syntax	<pre>show log <filename user <username>></pre>
Syntax (QFX Series and OCX Series)	<pre>show log filename <device-type (device-id device-alias)></pre>
Syntax (TX Matrix Router)	<pre>show log <all-lcc lcc number scc> <filename user <username>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <i>device-type (device-id device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.



NOTE: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options	<p>none—List all log files.</p> <p><all-lcc lcc number scc>—(Routing matrix only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).</p> <p>device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:</p>
----------------	--

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.



NOTE: If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

(device-id | device-alias)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include **username**, display logging information about the specified user.

Required Privilege Level trace

Related Documentation • *syslog (System)*

List of Sample Output [show log on page 156](#)
[show log filename on page 157](#)
[show log filename \(QFabric System\) on page 157](#)
[show log user on page 158](#)

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin     1056350 Sep 30 07:04 rpd.6
```

```
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin      19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr
192.0.2.21 nhop type local nhop 192.0.2.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
192.0.2.22 nhop type unicast nhop 192.0.2.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
```

```
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,  
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

```
user@host> show log user  
usera    mg2546                Thu Oct  1 19:37    still logged in  
usera    mg2529                Thu Oct  1 19:08 - 19:36 (00:28)  
usera    mg2518                Thu Oct  1 18:53 - 18:58 (00:04)  
root     mg1575                Wed Sep 30 18:39 - 18:41 (00:02)  
root     ttyp2      aaa.bbbb.com   Wed Sep 30 18:39 - 18:41 (00:02)  
userb    ttyp1      192.0.2.0    Wed Sep 30 01:03 - 01:22 (00:19)
```

show (ospf | ospf3) overview

List of Syntax	Syntax on page 159 Syntax (EX Series Switch and QFX Series) on page 159
Syntax	<pre>show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Database protection introduced in Junos 10.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display Open Shortest Path First (OSPF) overview information.
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf overview on page 161 show ospf overview (With Database Protection) on page 162 show ospf3 overview (With Database Protection) on page 162 show ospf overview extensive on page 162
Output Fields	<p>Table 10 on page 160 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear.</p>

Table 10: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: Current , Warning (threshold), and Allowed .	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: Current and Allowed .	All levels
Restart	Graceful restart capability: enabled or disabled .	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels

Table 10: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: enabled or disabled .	All levels
Trace options	OSPF-specific trace options.	extensive
Trace file	Name of the file to receive the output of the tracing operation.	extensive
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: Normal Stub , Not Stub , or Not so Stubby Stub .	All levels
Authentication Type	Type of authentication: None , Password , or MD5 .	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
Restart: Enabled
  Restart duration: 20 sec
  Restart grace period: 40 sec
  Helper mode: enabled
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 0
Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
Router ID: 10.255.112.218
Route table index: 0
LSA refresh time: 50 minutes
Traffic engineering
Restart: Enabled
  Restart duration: 180 sec
  Restart grace period: 210 sec
  Graceful restart helper mode: Enabled
  Restart-signaling helper mode: Enabled
Database protection state: Normal
Warning threshold: 70 percent
Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 1
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 70
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
Router ID: 10.255.112.128
Route table index: 0
LSA refresh time: 50 minutes
Database protection state: Normal
Warning threshold: 80 percent
Non self-generated LSAs: Current 3, Warning 8, Allowed 10
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 2
Area: 0.0.0.0
  Stub type: Not Stub
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 7
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
Router ID: 1.1.1.103
Route table index: 0
Full SPF runs: 13, SPF delay: 0.200000 sec
LSA refresh time: 50 minutes
```

```
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```


CHAPTER 11

Operational Commands for Graceful Switchover

- `show system switchover`
- `show task replication`

show system switchover

List of Syntax	Syntax on page 166
	Syntax (TX Matrix Router) on page 166
	Syntax (TX Matrix Plus Router) on page 166
	Syntax (MX Series Router) on page 166
Syntax	show system switchover
Syntax (TX Matrix Router)	show system switchover <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system switchover <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system switchover <all-members> <local> <member <i>member-id</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.
Description	Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.



NOTE: Issue the `show system switchover` command *only* on the backup Routing Engine. This command is *not* supported on the master Routing Engine, because the kernel-replication process daemon does not run on the master Routing Engine. This process runs only on the backup Routing Engine.

Beginning Junos OS Release 9.6, the `show system switchover` command has been deprecated on the master Routing Engine on all routers other than a TX Matrix (switch-card chassis) or a TX Matrix Plus (switch-fabric chassis) router.

However, in a routing matrix, if you issue the `show system switchover` command on the master Routing Engine of the TX Matrix router (or switch-card chassis), the CLI displays graceful switchover information for the master Routing Engine of the T640 routers (or line-card chassis) in the routing matrix. Likewise, if you issue the `show system switchover` command on the master Routing Engine of a TX Matrix Plus router (or switch-fabric chassis), the CLI displays output for the master Routing Engine of T1600 or T4000 routers in the routing matrix.

Options **all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix router and the T640 routers configured in the routing matrix. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix Plus router and the T1600 or T4000 routers configured in the routing matrix.

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all connected T1600 or T4000 LCCs.

Note that in this instance, packets get dropped. The LCCs perform GRES on their own chassis (GRES cannot be handled by one particular chassis for the entire router) and synchronization is not possible as the LCC plane bringup time varies for each LCC. Therefore, when there is traffic on these planes, there may be a traffic drop.

all-members—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on all members of the Virtual Chassis configuration.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for a specific T640 router connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for a specific router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display graceful Routing Engines switchover information for all Routing Engines on the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc—(TX Matrix router only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix Plus router.

Additional Information If you issue the **show system switchover** command on a TX Matrix backup Routing Engine, the command is broadcast to all the T640 backup Routing Engines that are connected to it.

Likewise, if you issue the **show system switchover** command on a TX Matrix Plus backup Routing Engine, the command is broadcast to all the T1600 or T4000 backup Routing Engines that are connected to it.

If you issue the **show system switchover** command on the active Routing Engine in the master router of an MX Series Virtual Chassis, the router displays a message that this command is not applicable on this member of the Virtual Chassis.

Required Privilege Level view

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output

[show system switchover \(Backup Routing Engine - Ready\) on page 169](#)
[show system switchover \(Backup Routing Engine - Not Ready\) on page 169](#)
[show system switchover \(MX Virtual Chassis\) on page 169](#)
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Master Ready on page 170](#)
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Master Not Ready on page 170](#)
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Backup Ready on page 170](#)
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Backup Not Ready on page 171](#)
[show system switchover all-icc \(Routing Matrix and Routing Matrix Plus\) on page 171](#)

Output Fields [Table 11 on page 168](#) describes the output fields for the **show system switchover** command. Output fields are listed in the approximate order in which they appear.

Table 11: show system switchover Output Fields

Field Name	Field Description
Graceful switchover	Display graceful Routing Engine switchover status: <ul style="list-style-type: none"> • On—Indicates graceful-switchover is specified for the routing-options configuration command. • Off—Indicates graceful-switchover is not specified for the routing-options configuration command.
Configuration database	State of the configuration database: <ul style="list-style-type: none"> • Ready—Configuration database has synchronized. • Synchronizing—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds. • Synchronize failed—Configuration database synchronize process failed.

Table 11: show system switchover Output Fields (*continued*)

Field Name	Field Description
Kernel database	<p>State of the kernel database:</p> <ul style="list-style-type: none"> • Ready—Kernel database has synchronized. This message implies that the system is ready for GRES. • Synchronizing—Kernel database is synchronizing. Displayed when there are updates within the last 5 seconds. • Version incompatible—The primary and standby Routing Engines are running incompatible kernel database versions. • Replication error—An error occurred when the state was replicated from the primary Routing Engine. Inspect Steady State for possible causes, or notify Juniper Networks customer support.
Peer state	<p>Routing Engine peer state:</p> <p>This field is displayed only when ksyncd is running in multichassis mode (LCC master).</p> <ul style="list-style-type: none"> • Steady State—Peer completed switchover transition. • Peer Connected—Peer in switchover transition.
Switchover	<p>Switchover status (output of master switch check command):</p> <ul style="list-style-type: none"> • Ready—Message for system being switchover ready. • error: Command aborted. Not ready for mastership switch, try after xxx secs.

Sample Output

show system switchover (Backup Routing Engine - Ready)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Ready
```

show system switchover (Backup Routing Engine - Not Ready)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

```
error: Command aborted. Not ready for mastership switch, try after 174 secs.
```

show system switchover (MX Virtual Chassis)

```
{master:member1-re1}
user@host> show system switchover
member0:
-----
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready
```

member1:

Command is not applicable on this member of the virtual-chassis

show system switchover (Routing Matrix and Routing Matrix Plus) - Master Ready

```
user@host> show system switchover
lcc0-re1:
```

Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Ready

lcc2-re0:

Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Ready

show system switchover (Routing Matrix and Routing Matrix Plus) - Master Not Ready

```
user@host> show system switchover
lcc0-re1:
```

Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Ready

lcc2-re1:

Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
error: Command aborted. Not ready for mastership switch, try after 228 secs.

show system switchover (Routing Matrix and Routing Matrix Plus) - Backup Ready

```
user@host> show system switchover
scc-re0:
```

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready

lcc0-re0:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready

lcc2-re1:

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready
```

show system switchover (Routing Matrix and Routing Matrix Plus) - Backup Not Ready

```
user@host> show system switchover
scc-re0:
-----
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
error: Command aborted. Not ready for mastership switch, try after 223 secs.

lcc0-re0:
-----
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready

lcc2-re1:
-----
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready
```

show system switchover all-lcc (Routing Matrix and Routing Matrix Plus)

```
user@host> show system switchover all-lcc

lcc0-re0:
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
lcc2-re0:
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

show task replication

Syntax	show task replication
Release Information	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.</p> <p>Support for logical systems introduced in Junos OS Release 13.3</p>
Description	Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show task replication (Issued on the Master Routing Engine) on page 172 show task replication (Issued on the Backup Routing Engine) on page 173
Output Fields	Table 12 on page 172 lists the output fields for the show task replication command. Output fields are listed in the approximate order in which they appear.

Table 12: show task replication Output Fields

Field Name	Field Description
Stateful replication	Displays whether or not graceful Routing Engine switchover is configured. The status can be Enabled or Disabled .
RE mode	Displays the Routing Engine on which the command is issued: Master , Backup , or Not applicable (when the router has only one Routing Engine).
Protocol	Protocols that are supported by nonstop active routing.
Synchronization Status	Nonstop active routing synchronization status for the supported protocols. States are NotStarted , InProgress , and Complete .

Sample Output

show task replication (Issued on the Master Routing Engine)

```

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

  Protocol              Synchronization Status
  OSPF                  NotStarted
  BGP                   Complete
  IS-IS                 NotStarted

```

LDP	Complete
PIM	Complete

show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication
Stateful Replication: Enabled
RE mode: Backup
```


CHAPTER 12

Operational Command for Nonstop Routing

- `show task replication`

show task replication

Syntax	show task replication
Release Information	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.</p> <p>Support for logical systems introduced in Junos OS Release 13.3</p>
Description	Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show task replication (Issued on the Master Routing Engine) on page 176 show task replication (Issued on the Backup Routing Engine) on page 177
Output Fields	Table 12 on page 172 lists the output fields for the show task replication command. Output fields are listed in the approximate order in which they appear.

Table 13: show task replication Output Fields

Field Name	Field Description
Stateful replication	Displays whether or not graceful Routing Engine switchover is configured. The status can be Enabled or Disabled .
RE mode	Displays the Routing Engine on which the command is issued: Master , Backup , or Not applicable (when the router has only one Routing Engine).
Protocol	Protocols that are supported by nonstop active routing.
Synchronization Status	Nonstop active routing synchronization status for the supported protocols. States are NotStarted , InProgress , and Complete .

Sample Output

show task replication (Issued on the Master Routing Engine)

```

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

  Protocol      Synchronization Status
  OSPF          NotStarted
  BGP           Complete
  IS-IS         NotStarted

```


LDP	Complete
PIM	Complete

show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication
Stateful Replication: Enabled
RE mode: Backup
```


CHAPTER 13

Operational Command for NSSU

- request system software nonstop-upgrade

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)*
- *Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*
- [Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade on page 79](#)
- *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 183](#)
[request system software nonstop-upgrade \(EX6200 Switch\) on page 185](#)
[request system software nonstop-upgrade reboot \(EX8200 Switch\) on page 186](#)
[request system software nonstop-upgrade no-old-master-upgrade \(EX8200 Switch\) on page 187](#)
[request system software nonstop-upgrade reboot \(EX8200 Virtual Chassis\) on page 187](#)

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

Sample Output

[request system software nonstop-upgrade \(EX4200 Virtual Chassis\)](#)

```
user@switch> request system software nonstop-upgrade
```

```
/var/tmp/jinstall-ex-4200-12.1R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Installing image on other FPC's along with the backup

Checking pending install on fpc1
Pushing bundle to fpc1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc1

Checking pending install on fpc2
Pushing bundle to fpc2
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc2

Checking pending install on fpc3
Pushing bundle to fpc3
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc3

Checking pending install on fpc4
Pushing bundle to fpc4
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc4

Checking pending install on fpc5
Pushing bundle to fpc5
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc5

Checking pending install on fpc6
Pushing bundle to fpc6
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc6

Checking pending install on fpc7
Pushing bundle to fpc7
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc7
Backup upgrade done
Rebooting Backup RE

Rebooting fpc1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
```



```

Checking In-Service-Upgrade status
Item           Status           Reason
FPC 0          Online
FPC 1          Online
FPC 2          Online (ISSU)
FPC 3          Online (ISSU)
FPC 4          Online (ISSU)
FPC 5          Online (ISSU)
FPC 6          Online (ISSU)
FPC 7          Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE

*** FINAL System shutdown message from root@switch ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 9336]

```

request system software nonstop-upgrade (EX6200 Switch)

```

{master}
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-6200-12.2R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re0
NOTICE: Validating configuration against
jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re0
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
Item           Status           Reason
FPC 0          Online (ISSU)
FPC 1          Online (ISSU)
FPC 2          Online (ISSU)
FPC 3          Online (ISSU)
FPC 4          Online
FPC 5          Online

```

```

FPC 6      Online (ISSU)
FPC 7      Online (ISSU)
FPC 8      Online (ISSU)
FPC 9      Online (ISSU)
Going to install image on master
NOTICE: Validating configuration against
jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE
Trying to relinquish mastership before rebooting...
Resolving mastership...
Complete. The other routing engine becomes the master.

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY

```

request system software nonstop-upgrade reboot (EX8200 Switch)

```

{master}
user@switch> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item      Status      Reason
  FPC 0     Online (ISSU)
  FPC 2     Offline      Offlined by CLI command
  FPC 3     Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
[pid 2635]

```

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

request system software nonstop-upgrade no-old-master-upgrade (EX8200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade no-old-master-upgrade
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 3         Offline
  FPC 4         Online (ISSU)
  FPC 5         Online (ISSU)
  FPC 6         Online (ISSU)
  FPC 7         Online (ISSU)
  Offlined by CLI command
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
```


CHAPTER 14

Operational Commands for VRRP

- `show vrrp`

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)
List of Sample Output	<p>show vrrp on page 197</p> <p>show vrrp brief on page 197</p> <p>show vrrp detail (IPv6) on page 197</p> <p>show vrrp detail (Route Track) on page 198</p> <p>show vrrp extensive on page 198</p> <p>show vrrp interface on page 200</p> <p>show vrrp summary on page 200</p> <p>show vrrp track detail on page 200</p> <p>show vrrp track summary on page 201</p>
Output Fields	Table 14 on page 192 lists the output fields for the show vrrp command. Output fields are listed in the approximate order in which they appear.

Table 14: 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 14: 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 14: 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 14: 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 14: 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	<p>Errored statistics for the VRRP group:</p> <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	<p>State transition statistics for the VRRP group:</p> <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 14: 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
ge2001:db8::12:1:1:99
ge-0/0/2.131   up         1      master    A 0.364 1c1
ge2001:db8::13:1:1:1
ge2001:db8:0:1:13:1:1:99
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 197](#).

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

