

High Availability Feature Guide for QFabric Systems

Release

14.1X53-D15



Modified: 2015-08-26

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

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

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

High Availability Feature Guide for QFabric Systems
14.1X53-D15
Copyright © 2015, 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 Graceful Restart for QFabric Systems	4
	Enabling Graceful Restart	4
	Configuring Graceful Restart Options for BGP	5
	Configuring Graceful Restart Options for OSPF and OSPFv3	6
	Tracking Graceful Restart Events	7
	Configuring Routing Protocols Graceful Restart	8
	Enabling Graceful Restart	8
	Configuring Graceful Restart Options for BGP	9
	Configuring Graceful Restart Options for ES-IS	10
	Configuring Graceful Restart Options for IS-IS	10
	Configuring Graceful Restart Options for OSPF and OSPFv3	11
	Configuring Graceful Restart Options for RIP and RIPng	12
	Configuring Graceful Restart Options for PIM Sparse Mode	13
	Tracking Graceful Restart Events	14
Part 2	Graceful Routing Engine Switchover	
Chapter 2	Using GRES	17
	Understanding Graceful Routing Engine Switchover	17
	Graceful Routing Engine Switchover Concepts	17
	Effects of a Routing Engine Switchover	21
	Graceful Routing Engine Switchover System Requirements	23
	Graceful Routing Engine Switchover Platform Support	23
	Graceful Routing Engine Switchover Feature Support	24
	Graceful Routing Engine Switchover DPC Support	25

	Graceful Routing Engine Switchover and Subscriber Access	25
	Graceful Routing Engine Switchover PIC Support	25
	Configuring Graceful Routing Engine Switchover	26
	Enabling Graceful Routing Engine Switchover	26
	Configuring Graceful Routing Engine Switchover with Graceful Restart	26
	Synchronizing the Routing Engine Configuration	27
	Verifying Graceful Routing Engine Switchover Operation	27
	Resetting Local Statistics	28
Part 3	Nonstop Software Upgrade	
Chapter 3	Using NSSU	31
	Understanding Nonstop Software Upgrade for QFabric Systems	31
	Performing a Nonstop Software Upgrade on the QFabric System	36
	Backing Up the Current Configuration Files	37
	Downloading Software Files Using a Browser	37
	Retrieving Software Files for Download	39
	Performing a Nonstop Software Upgrade for Director Devices in a Director Group	39
	Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components	39
	(Optional) Creating Upgrade Groups for Node Groups	40
	Performing a Nonstop Software Upgrade on a Node Group	40
	Verifying Nonstop Software Upgrade for QFabric Systems	41
	Verifying a Director Group Nonstop Software Upgrade	41
	Verifying a Fabric Nonstop Software Upgrade	54
	Verifying a Redundant Server Node Group Nonstop Software Upgrade	55
	Verifying a Network Node Group Nonstop Software Upgrade	59
Part 4	Virtual Router Redundancy Protocol	
Chapter 4	Using VRRP	63
	Understanding VRRP	63
	Overview of VRRP	63
	Sample VRRP Topology	64
	Understanding VRRP Between QFabric Systems	66
	VRRP Differences on QFabric Systems	66
	Configuration Details	66
	Example: Configuring VRRP for Load Sharing	69
	Configuring Basic VRRP Support	74
	Configuring VRRP Authentication (IPv4 Only)	75
	Configuring the Startup Period for VRRP Operations	76
	Configuring the Advertisement Interval for the VRRP Master	76
	Modifying the Advertisement Interval in Seconds	77
	Modifying the Advertisement Interval in Milliseconds	77
	Configuring VRRP Preemption and Hold Time	77
	Configuring VRRP Preemption	78
	Configuring the Preemption Hold Time	78
	Overriding the Hold Time	78

	Configuring a Route to Be Tracked	79
	Configuring a Logical Interface to Be Tracked	79
	Configuring a Backup to Accept Packets Destined for the Virtual IP Address	81
	Configuring Passive ARP Learning for VRRP Backups	82
	Configuring the Silent Period	82
	Configuring Inheritance for a VRRP Group	83
Part 5	Configuration Statements and Operational Commands	
Chapter 5	Configuration Statements for Graceful Restart	87
	disable	88
	disable (BGP Graceful Restart)	89
	graceful-restart (Enabling Globally)	90
	graceful-restart (Protocols BGP)	92
	graceful-restart (Protocols OSPF)	93
	helper-disable (OSPF)	95
	no-strict-lsa-checking	96
	notify-duration	97
	redundancy (Graceful Switchover)	98
	restart-duration	99
	restart-time (BGP Graceful Restart)	100
	stale-routes-time	101
Chapter 6	Configuration Statement for Graceful Switchover	103
	graceful-switchover	103
	redundancy (Graceful Switchover)	104
Chapter 7	Configuration Statements for VRRP	105
	accept-data	106
	advertise-interval	107
	asymmetric-hold-time	108
	authentication-key	109
	authentication-type	110
	bandwidth-threshold	111
	failover-delay	112
	fast-interval	113
	hold-time (VRRP)	114
	interface (VRRP Group)	115
	preempt (VRRP)	116
	priority (Protocols VRRP)	117
	priority-cost (VRRP)	118
	priority-hold-time	119
	route (Interfaces)	120
	startup-silent-period	121
	traceoptions	122
	track (VRRP)	124
	virtual-address	125
	vrrp-group	126

Chapter 8	Operational Commands for Graceful Restart	129
	Verifying Graceful Restart Operation	129
	Graceful Restart Operational Mode Commands	129
	Verifying BGP Graceful Restart	130
	Verifying IS-IS and OSPF Graceful Restart	130
	Verifying CCC and TCC Graceful Restart	131
	show bgp neighbor	132
	show log	146
	show (ospf ospf3) overview	149
Chapter 9	Operational Command for Graceful Switchover	155
	show system switchover	156
	show task replication	160
Chapter 10	Operational Commands for Nonstop Software Upgrade	163
	request system software nonstop-upgrade	164
	show chassis nonstop-upgrade node-group	170
Chapter 11	Operational Commands for VRRP	171
	show vrrp	172

List of Figures

Part 2	Graceful Routing Engine Switchover	
Chapter 2	Using GRES	17
	Figure 1: Preparing for a Graceful Routing Engine Switchover	19
	Figure 2: Graceful Routing Engine Switchover Process	20
Part 4	Virtual Router Redundancy Protocol	
Chapter 4	Using VRRP	63
	Figure 3: Basic VRRP Topology	65
	Figure 4: 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	17
	Table 3: Effects of a Routing Engine Switchover	22
	Table 4: Graceful Routing Engine Switchover Feature Support	24
Part 4	Virtual Router Redundancy Protocol	
Chapter 4	Using VRRP	63
	Table 5: RVIs on QFabric systems in example VRRP configuration	67
	Table 6: Sample VRRP configuration each RVI	67
	Table 7: Interfaces on QFabric system A. All interfaces are members of VLAN 100.	68
	Table 8: Interfaces on QFabric system B. All interfaces are members of VLAN 100 (same as on QFabric system A).	68
	Table 9: Settings for VRRP Load-Sharing Example	70
	Table 10: Interface State and Priority Cost Usage	81
Part 5	Configuration Statements and Operational Commands	
Chapter 8	Operational Commands for Graceful Restart	129
	Table 11: show bgp neighbor Output Fields	133
	Table 12: show ospf overview Output Fields	150
Chapter 9	Operational Command for Graceful Switchover	155
	Table 13: show system switchover Output Fields	158
	Table 14: show task replication Output Fields	160
Chapter 10	Operational Commands for Nonstop Software Upgrade	163
	Table 15: show chassis nonstop-upgrade node-group Output Fields	170
Chapter 11	Operational Commands for VRRP	171
	Table 16: show vrrp Output Fields	172

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:

- QFabric System

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

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

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

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

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

Documentation Feedback

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

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation 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 <https://www.juniper.net/cgi-bin/docbugreport/>.
- 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 Graceful Restart for QFabric Systems on page 4](#)
- [Configuring Routing Protocols Graceful Restart on page 8](#)

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 on page 4](#)

Configuring Graceful Restart for QFabric Systems

When you configure graceful restart in the QFabric CLI, the QFabric system applies the configuration to the network Node group to participate in graceful restart operations with devices external to the QFabric system. Such configuration preserves routing table state and helps neighboring routing devices to resume routing operations more quickly after a system restart. This also enables the network Node group to resume routing operations rapidly if there is a restart in the QFabric system (such as a software upgrade). As a result, we recommend enabling graceful restart for routing protocols in the QFabric CLI.



NOTE: The QFabric system also uses graceful restart internally within the fabric to facilitate interfabric resiliency and recovery. This internal feature is enabled by default with no configuration required.

- [Enabling Graceful Restart on page 4](#)
- [Configuring Graceful Restart Options for BGP on page 5](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 6](#)
- [Tracking Graceful Restart Events on page 7](#)

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 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 | ospfv3) 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 reenabling 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 7](#).



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

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](#)
 - [Verifying Graceful Restart Operation on page 129](#)

Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- [Enabling Graceful Restart on page 8](#)
- [Configuring Graceful Restart Options for BGP on page 9](#)
- [Configuring Graceful Restart Options for ES-IS on page 10](#)
- [Configuring Graceful Restart Options for IS-IS on page 10](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 11](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 12](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 13](#)
- [Tracking Graceful Restart Events on page 14](#)

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. See *Example: Configuring IS-IS for GRES with Graceful Restart* for more information.



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 7](#).

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 | ospfv3) 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 reenable 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 | ospfv3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 7](#).



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 129](#)
- *Configuring Graceful Restart*
- *Example: Configuring IS-IS for GRES with Graceful Restart*

PART 2

Graceful Routing Engine Switchover

- [Using GRES on page 17](#)

CHAPTER 2

Using GRES

- [Understanding Graceful Routing Engine Switchover on page 17](#)
- [Graceful Routing Engine Switchover System Requirements on page 23](#)
- [Configuring Graceful Routing Engine Switchover on page 26](#)
- [Resetting Local Statistics on page 28](#)

Understanding Graceful Routing Engine Switchover

This topic contains the following sections:

- [Graceful Routing Engine Switchover Concepts on page 17](#)
- [Effects of a Routing Engine Switchover on page 21](#)

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*.

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: 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.

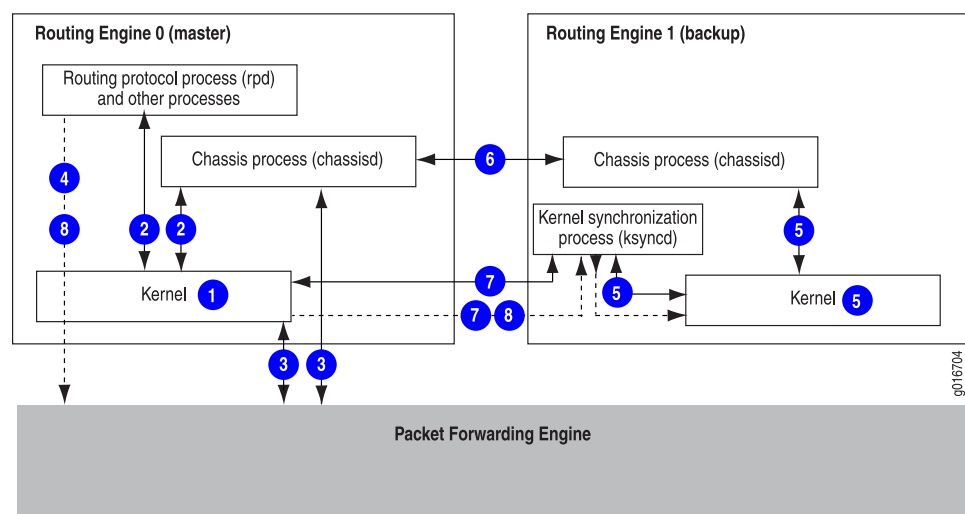


NOTE:

- 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 19 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:

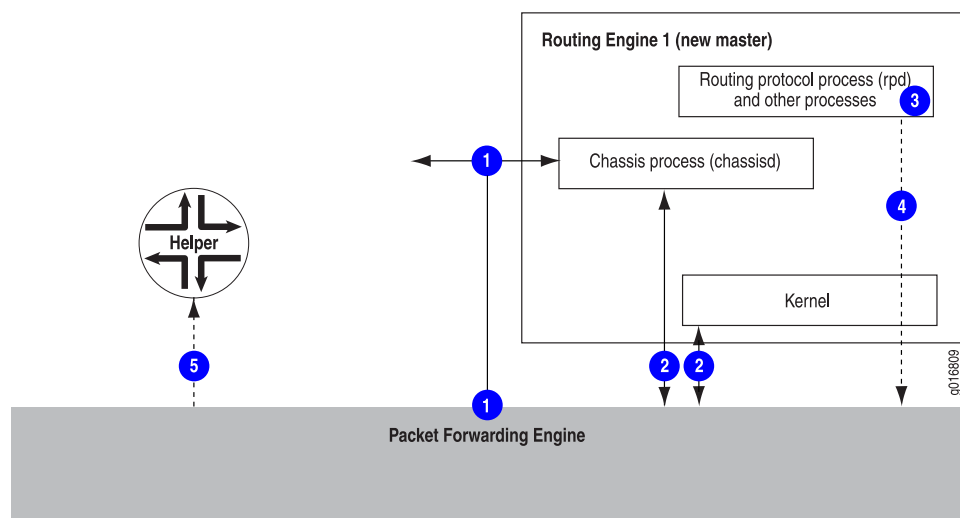
- The **request chassis routing-engine master switch check** command from the master Routing Engine
- The **show system switchover** command from the Backup Routing Engine

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 20 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 22 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.
GRES <i>and</i> 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 23](#)
- [Configuring Graceful Routing Engine Switchover on page 26](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\)](#)
- [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 23](#)
- [Graceful Routing Engine Switchover Feature Support on page 24](#)
- [Graceful Routing Engine Switchover DPC Support on page 25](#)
- [Graceful Routing Engine Switchover and Subscriber Access on page 25](#)
- [Graceful Routing Engine Switchover PIC Support on page 25](#)

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 24](#).

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)
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 is the default) and that VRRP is not running on logical interfaces like aggregate-interface (ae) or integrated-routing-bridging-interfaces (irb).

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 23](#). 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 17](#)
- [Configuring Graceful Routing Engine Switchover on page 26](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\)](#)
- [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 26](#)
- [Configuring Graceful Routing Engine Switchover with Graceful Restart on page 26](#)
- [Synchronizing the Routing Engine Configuration on page 27](#)
- [Verifying Graceful Routing Engine Switchover Operation on page 27](#)

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.

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 , 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 17](#)
- [Graceful Routing Engine Switchover System Requirements on page 23](#)
- [Requirements for Routers with a Backup Router Configuration](#)
- [Resetting Local Statistics on page 28](#)
- [graceful-switchover](#)
- [graceful-switchover on page 103](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)
- [hold-time](#)

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 17](#)
- [Configuring Graceful Routing Engine Switchover on page 26](#)

PART 3

Nonstop Software Upgrade

- [Using NSSU on page 31](#)

CHAPTER 3

Using NSSU

- [Understanding Nonstop Software Upgrade for QFabric Systems on page 31](#)
- [Performing a Nonstop Software Upgrade on the QFabric System on page 36](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 41](#)

Understanding Nonstop Software Upgrade for QFabric Systems

The framework that underlies a nonstop software upgrade in a QFabric system enables you to upgrade the system in a step-by-step manner and minimize the impact to the continuous operation of the system. This topic explains how a nonstop software upgrade works in a QFabric system, the steps that are involved, and the procedures that you need to implement to experience the benefits of this style of software upgrade.

Nonstop software upgrade enables some QFabric system components to continue operating while similar components in the system are being upgraded. In general, the QFabric system upgrades redundant components in stages so that some components remain operational and continue forwarding traffic while their equivalent counterparts upgrade to a new version of software.



TIP: Use the following guidelines to decide when to implement a nonstop software upgrade:

- If you need to upgrade all components of the system in the shortest amount of time (approximately one hour) and you do not need to retain the forwarding resiliency of the data plane, issue the **request system software add component all** command to perform a standard software upgrade. All components of the QFabric system upgrade simultaneously and expediently, but this type of upgrade does not provide resiliency or switchover capabilities.
- If you need to minimize service impact, preserve the forwarding operations of the data plane during the upgrade, and are willing to take the extra time required for component switchovers (in many cases, several hours), issue the three nonstop software upgrade commands (**request system software nonstop-upgrade (director-group | fabric | node-group)**) described in this topic in the correct order.



NOTE:

- Before you begin a nonstop software upgrade, issue the **request system software download** command to copy the software to the QFabric system.
 - Each of the 3 nonstop software upgrade steps must be considered parts of the whole process. You must complete all 3 steps of a nonstop software upgrade in the correct order to ensure the proper operation of the QFabric system.
 - Open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a second session to verify that the QFabric system components respond to operational mode commands as expected. For more information on verification of the upgrade, see [“Verifying Nonstop Software Upgrade for QFabric Systems” on page 41](#).
 - Issue the **show fabric administration inventory** command to verify that all upgraded components are operational at the end of a step before beginning the next step.
 - Once you start the nonstop software upgrade process, we strongly recommend that you complete all 3 steps within 12 hours.
-

The three steps to a successful nonstop software upgrade must be performed in the following order:

- Director group—The first step upgrades the Director devices, the fabric manager Routing Engine, and the diagnostic Routing Engine. To perform the first step, issue the **request system software nonstop-upgrade director-group** command. The key actions that occur during a Director group upgrade are:
 1. Connecting to the QFabric system by way of an SSH connection. This action establishes a load-balanced CLI session on one of the Director devices in the Director group.
 2. The QFabric system downloads and installs the new software in both Director devices.
 3. The Director device hosting the CLI session becomes the master for all QFabric system processes running on the Director group, such as the fabric manager and network Node group Routing Engines.
 4. The QFabric system installs the new software for the backup fabric manager Routing Engine on the backup Director device.
 5. The backup Director device reboots to activate the new software.
 6. The master Director device begins a 15 minute sequence that includes a temporary suspension of QFabric services and a QFabric database transfer. You cannot issue operational mode commands in the QFabric CLI during this period.
 7. The QFabric system installs the new software for the fabric manager and diagnostic Routing Engines on the Director group master.

8. The QFabric system switches mastership of all QFabric processes from the master Director device to the backup Director device.
9. The master Director device reboots to activate the new software.
10. The CLI session terminates, and logging back in to the QFabric system with a new SSH connection establishes the session on the new master Director device (the original backup).
11. The previous master Director device resumes operation as a backup and the associated processes (such as the fabric manager and network Node group Routing Engines) become backup as well. The fabric control Routing Engine associated with this Director device returns to active status.



NOTE: After the Director group nonstop software upgrade completes, any Interconnect device or Node device that reboots will automatically download the new software, install it, and reboot again. As a result, try not to restart any QFabric system devices before you complete the rest of the nonstop software upgrade steps.



TIP:

- To enable BGP and OSPF to continue operating on the network Node group during a Director group nonstop service upgrade, we recommend that you configure graceful restart for these routing protocols. For more information on graceful restart, see [“Configuring Graceful Restart for QFabric Systems” on page 4](#).
 - Wait 15 minutes after the second Director device returns to service and hosts Routing Engine processes before proceeding to step 2—the fabric upgrade. You can verify the operational status of both Director devices by issuing the `show fabric administration inventory director-group status` command. Also, issue the `show fabric administration inventory infrastructure` command to verify when the Routing Engine processes become load balanced (typically, there will be three to four Routing Engines running on each Director device).
-
- Fabric—The second step upgrades the Interconnect devices and the fabric control Routing Engines. To perform the second step, issue the **request system software nonstop-upgrade fabric** command. The key actions that occur during a fabric upgrade are:
 1. The QFabric system downloads, validates, and installs the new software in all Interconnect devices and fabric control Routing Engines (FC-0 and FC-1).
 2. One fabric control Routing Engine reboots and comes back online.
 3. The other fabric control Routing Engine reboots and comes back online.

4. The first Interconnect device reboots, comes back online, and resumes the forwarding of traffic.
5. Subsequent Interconnect devices reboot one at a time, come back online, and return to service.

**NOTE:**

- If the software does not load properly on any one of the fabric components, all components revert back to the original software version.
- If one of the components in a fabric upgrade does not reboot successfully, issue the **request system reboot fabric** command to reattempt the rebooting process for this fabric component and activate the new software.

- **Node group**—The third and final step upgrades Node groups. You can choose to upgrade a network Node group, a redundant server Node group, or individual server Node groups. You can upgrade the Node groups one at a time or in groups (known as upgrade groups). However, you must upgrade all Node groups in your QFabric system before you can complete the nonstop software upgrade process. To perform the third step, issue the **request system software nonstop-upgrade node-group** command.

The key actions that occur during a network Node group upgrade are:

1. The QFabric system copies the new software to each Node device one at a time.
2. The QFabric system validates and then installs the new software in all Node devices simultaneously.
3. The system copies the software to the network Node group Routing Engines.
4. The QFabric system validates and then installs the software in the network Node group Routing Engines one at a time -- first the backup, then the master.
5. The backup network Node group Routing Engine reboots and comes back online.
6. The supporting Node devices reboot and come back online one at a time.



NOTE: To reduce the total upgrade duration, configure an upgrade group. All Node devices within the upgrade group reboot at the same time.

7. The master network Node group Routing Engine relinquishes mastership to the backup, reboots, and comes back online.

The key actions that occur during a redundant server Node group upgrade are:

1. The QFabric system copies the new software to the backup Node device, then the master Node device.
2. The QFabric system validates and then installs the new software on the backup Node device, then the master Node device.

3. The backup Node device reboots, comes back online, and becomes the master Node device.
4. The previous master Node device reboots and comes back online as a backup Node device.



NOTE: For redundant server Node groups, both Node devices must be online before the upgrade will proceed. If one of the devices is no longer available, remove the Node device from the Node group configuration before you issue the nonstop software upgrade command.

The key actions that occur during a server Node group upgrade for a Node group that contains one member are:

1. The Node device downloads the software package and validates the software.
2. The Node device installs the software and reboots.



NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade.

Related Documentation

- [Performing a Nonstop Software Upgrade on the QFabric System on page 36](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 41](#)
- [request system software nonstop-upgrade on page 164](#)
- [request system software add](#)
- [Configuring Graceful Restart for QFabric Systems on page 4](#)

Performing a Nonstop Software Upgrade on the QFabric System



NOTE: Before you can perform a nonstop software upgrade to Junos OS Release 13.1X50-D10, you must have Junos OS Release 12.2X50-D42 or later installed. You cannot perform a nonstop software upgrade with Junos OS Release 12.2X50-D41 or earlier. Contact the Juniper Technical Assistance Center for information on how to download Junos OS Release 12.2X50-D42. Performing a standard software upgrade (that is, issuing the `request system software add component all` command) does not require that you upgrade to an intermediate Junos OS software release.

To perform a nonstop software upgrade to Junos OS Release 13.1X50-D10:

1. First perform a nonstop software upgrade to Junos OS Release 12.2X50-D42.
2. Then perform a nonstop software upgrade to Junos OS Release 13.1X50-D10.

Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. This feature introduces several high availability improvements to the QFabric system software upgrade process, including:

- Upgrading members of a Director group or Node group one at a time so that one device in the group is always operational
- Switching mastership of Routing Engine processes to the backup Director device before upgrading the master Director device
- Rebooting Interconnect devices and fabric control Routing Engines one at a time, so that one Interconnect device or one fabric control Routing Engine is always operational
- Switching mastership of a Node group to the backup Node device before upgrading the master Node device
- Specifying an upgrade group if you want all Node devices in a Node group to be upgraded in parallel (which shortens the time of the upgrade)
- Rebooting devices automatically as part of the nonstop upgrade process

When performing a nonstop upgrade, start with the Director group upgrade, then issue the fabric upgrade, and end with the Node group upgrades.



NOTE: Because there is no redundancy for Node groups containing a single Node device, traffic loss occurs when the device reboots during the upgrade. For node-groups defined with two node-devices, both must be online in order for upgrade to succeed.



NOTE: Before you install the software, we recommend that you back up your current configuration files by issuing the `request system software configuration-backup` command.



NOTE: Before you can perform a nonstop software upgrade in your QFabric system, you must first upgrade your system to Junos OS Release 12.2 by using a conventional upgrade method such as issuing the `request system software add component all` command.

This topic describes the following tasks:

- [Backing Up the Current Configuration Files on page 37](#)
- [Downloading Software Files Using a Browser on page 37](#)
- [Retrieving Software Files for Download on page 39](#)
- [Performing a Nonstop Software Upgrade for Director Devices in a Director Group on page 39](#)
- [Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components on page 39](#)
- [\(Optional\) Creating Upgrade Groups for Node Groups on page 40](#)
- [Performing a Nonstop Software Upgrade on a Node Group on page 40](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

```
user@qfabric> request system software configuration-backup path
```

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@qfabric> request system software configuration-backup/media/USB/
```

Downloading Software Files Using a Browser



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the **Switching** box, click **Junos OS Platforms**.

4. In the **QFX Series** section, click the name of the platform for which you want to download software.
5. Click the **Software** tab and select the release number from the **Release** drop-down list.
6. Select the complete install package you want to download in the **QFabric System Install Package** section:
 - If you want to upgrade the entire QFabric system, select **QFabric System - Complete Install Package**.
 - If you want to upgrade either a single Node or Interconnect device for recovery purposes, select **Node and Interconnect Device Install Package**. For information on how to perform a recovery installation on either a Node or Interconnect device, see *Performing a Recovery Installation*.

A login screen appears.

7. Enter your user ID and password and press **Enter**.
8. Read the End User License Agreement, select the **I agree** option button, and then click **Proceed**.
9. Save the **jinstall-qfabric-version.rpm** file on your computer.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software:

```
user@qfabric> request system software download /path/package-name
```

For example:

```
user@qfabric> request system software download
ftp://server/files/jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Director Devices in a Director Group



NOTE: If you reboot any Node groups or Interconnect devices after you perform a nonstop upgrade on the Director group, these devices are upgraded to the same version of software that is running on the Director group.

To upgrade the software on the Director devices in a Director group:

- Issue the **request system software nonstop-upgrade director-group package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm
```

Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components

Before you perform a nonstop upgrade on the Interconnect devices and other fabric-related components, verify that both Director devices in the Director group are online. Both Director devices must be online before you attempt to perform a nonstop upgrade. To do verify that both Director devices are online, issue the **show fabric administration inventory director-group status** command.

To install the software on the Interconnect device and other components in the fabric:

- Issue the **request system software nonstop-upgrade fabric package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade fabric
jinstall-qfabric-12.2X50-D10.3.rpm
```

(Optional) Creating Upgrade Groups for Node Groups

Upgrade groups enable two or more Node devices in a Node group, or an entire Node group, to be rebooted at the same time. If you do not create an upgrade group, the Node devices are upgraded one at a time. Before performing a nonstop upgrade on a Node group, create an upgrade group and include the devices you want to reboot at the same time.



NOTE: If you add Node devices that have links to the same link aggregation group (LAG), there might be traffic loss.

- Create the upgrade group by issuing the **set chassis node-group *node-group-name* nssu upgrade-group *upgrade-group-name* node-devices** command at the [edit chassis] hierarchy.

For example:

```
user@qfabric# set chassis node-group nodegroup1 nssu upgrade-group upgrade1 node-devices
[ node1 node2 ]
```

Performing a Nonstop Software Upgrade on a Node Group

When you perform a nonstop software upgrade on a network Node group, the Node devices in the network Node group are upgraded in a serial fashion except when upgrade groups are configured. If you perform a nonstop upgrade on a redundant server Node group, both Node devices must be online for a successful upgrade. If one of the Node devices is no longer available, remove it from the configuration before you perform the nonstop software upgrade. If you perform a nonstop upgrade on a Node group with only one Node device, traffic loss occurs while the Node device is rebooting.



NOTE: You can upgrade multiple Node groups with this command. However, if more than one Node group is specified, there may be traffic loss depending on the topology of the network.

To install software on a Node group:

- Issue the **request system software nonstop-upgrade node-group *node-group-name* *package-name*** command.

To perform a nonstop upgrade on one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group nodegroup1
jinstall-qfabric-12.2X50-D10.3.rpm
```

To perform a nonstop upgrade on more than one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group [nodegroup1
nodegroup2 nodegroup3] jinstall-qfabric-12.2X50-D10.3.rpm
```

**Related
Documentation**

- [Configuring Graceful Restart for QFabric Systems on page 4](#)

- [Understanding Nonstop Software Upgrade for QFabric Systems on page 31](#)
- [Verifying Nonstop Software Upgrade for QFabric Systems on page 41](#)
- [request system software nonstop-upgrade on page 164](#)

Verifying Nonstop Software Upgrade for QFabric Systems

This topic discusses how you can monitor the progress of each of the three steps in a nonstop software upgrade. By identifying the key actions and events that define this process, you can track the status of the upgrade with confidence.



TIP: When performing a nonstop software upgrade, open two SSH sessions to the QFabric CLI. Use one session to monitor the upgrade itself and use a second session to verify that the QFabric system components respond to operational mode commands as expected.

- [Verifying a Director Group Nonstop Software Upgrade on page 41](#)
- [Verifying a Fabric Nonstop Software Upgrade on page 54](#)
- [Verifying a Redundant Server Node Group Nonstop Software Upgrade on page 55](#)
- [Verifying a Network Node Group Nonstop Software Upgrade on page 59](#)

Verifying a Director Group Nonstop Software Upgrade

Purpose During the Director group portion of a nonstop software upgrade, you should expect to see the Director device that hosts the CLI session selected as the master device. When mastership of all processes moves to the master, the QFabric system upgrades the backup Director device and this Director device reboots. After the backup Director device comes back online, the master Director device suspends CLI operations for 15 minutes, upgrades itself, and reboots. At this point, the backup becomes the new master Director device and you can issue CLI operational commands. Finally, the former master comes back online as a backup and both devices are operational once again. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

Action 1. In one SSH session to the QFabric CLI, verify the current status of the QFabric system by issuing the **show fabric administration inventory**, **show fabric administration inventory director-group status**, and **show fabric session-host** commands. In this case, Director device DG0 is the master device but DG1 hosts the CLI session.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
	P1507-C	Connected	
RSNG		Connected	Configured
	P1550-C	Connected	

P1571-C	Connected	
Interconnect device IC-F4912	Connected	Configured
F4912/RE0	Connected	
Fabric manager FM-0	Connected	Configured
Fabric control FC-0	Connected	Configured
FC-1	Connected	Configured
Diagnostic routing engine DRE-0	Connected	Configured

```
session1@qfabric> show fabric administration inventory director-group status
```

```
Director Group Status Tue Jun 5 15:11:26 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	8%	17363152k	4	3 days, 20:55 hrs
dg1	online	backup	10.49.215.39	6%	20157440k	3	3 days, 20:55 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	5.4G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

```

Director Group Processes
-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online      master

FTP Server                  online
Syslog                     online
Distributed Management      online
SNMP Trap Forwarder        online
SNMP Process                online
Platform Management         online

Interface Link Status
-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]               up
CP Link [0/1]               up
CP Link [0/0]               up
CP Link [1/2]               down
CP Link [1/1]               down
CP Link [1/0]               down
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032 online  backup

Director Group Managed Services
-----
Shared File System         online
Network File System        online
Virtual Machine Server     online
Load Balancer/DHCP         online

Hard Drive Status
-----
Volume ID:8                optimal
Physical ID:1               online
Physical ID:0               online
SCSI ID:1                   100%
SCSI ID:0                   100%

Size  Used Avail Used% Mounted on
-----
423G  5.5G 395G  2%  /
99M   16M  79M  17% /boot
93G   7.3G 86G   8%  /pbdata

Director Group Processes
-----
Director Group Manager      online

```

```

Partition Manager          online
Software Mirroring         online
Shared File System master  online
Secure Shell Process       online
Network File System        online
DHCP Server master        online    backup

```

```

FTP Server                online
Syslog                   online
Distributed Management    online
SNMP Trap Forwarder      online
SNMP Process             online
Platform Management      online

```

Interface Link Status

```

-----
Management Interface      up
Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG            up
CP Link [0/3]             up
CP Link [1/3]             up

```

```

session1@qfabric> show fabric session-host
Identifier: 0281052011000032

```

- In a second SSH session to the QFabric CLI, issue the request for the Director group nonstop software upgrade.

```

root@qfabric> request system software nonstop-upgrade director-group
install-qfabric-12.2X50-D10.3.rpm

```

- If the CLI session is being hosted by the master Director device, skip to step 4. However, if the CLI session is hosted by the backup Director device, the Director group mastership switches to the backup device after you issue the nonstop software upgrade command. In this example, mastership switches to Director device DG1.

```

session1@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun 5 15:12:20 UTC 2012

```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	backup	10.49.215.38	8%	31905924k	0	3 days, 21:16 hrs
dg1	online	master	10.49.215.39	6%	18010368k	3	3 days, 21:16 hrs

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	backup

Director Group Managed Services

```

-----
Shared File System        offline
Network File System       offline
Virtual Machine Server    offline

```

```

Load Balancer/DHCP                offline

Hard Drive Status
-----
Volume ID:4                      optimal
Physical ID:1                    online
Physical ID:0                    online
SCSI ID:1                        100%
SCSI ID:0                        100%

Size  Used Avail Used% Mounted on
----  -
423G  5.4G 395G  2%   /
99M   16M  79M   17%  /boot

Director Group Processes
-----
Director Group Manager           online
Partition Manager                online
Software Mirroring               online
Shared File System master        online
Secure Shell Process             online
Network File System              offline
DHCP Server master               offline    backup

FTP Server                       online
Syslog                          online
Distributed Management           offline
SNMP Trap Forwarder              offline
SNMP Process                     offline
Platform Management              online

Interface Link Status
-----
Management Interface             up
Control Plane Bridge             up
Control Plane LAG                up
CP Link [0/2]                    up
CP Link [0/1]                    up
CP Link [0/0]                    up
CP Link [1/2]                    down
CP Link [1/1]                    down
CP Link [1/0]                    down
Crossover LAG                    up
CP Link [0/3]                    up
CP Link [1/3]                    up

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032 online  master

Master Services
-----
Database Server              online
Load Balancer Director        online
QFabric Partition Address     online

Director Group Managed Services
-----
Shared File System            online

```

```

Network File System          online
Virtual Machine Server      online
Load Balancer/DHCP          online

```

Hard Drive Status

```

-----
Volume ID:8                  optimal
Physical ID:1                online
Physical ID:0                online
SCSI ID:1                    100%
SCSI ID:0                    100%

```

```

Size  Used Avail Used% Mounted on
-----
423G  6.0G 395G   2%   /
99M   16M  79M   17%  /boot
93G   7.3G 86G    8%  /pbdata

```

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager           online
Software Mirroring           online
Shared File System master   online
Secure Shell Process         online
Network File System          online
DHCP Server master           online      master

FTP Server                   online
Syslog                       online
Distributed Management        online
SNMP Trap Forwarder          online
SNMP Process                  online
Platform Management          online

```

Interface Link Status

```

-----
Management Interface        up
Control Plane Bridge         up
Control Plane LAG            up
CP Link [0/2]                up
CP Link [0/1]                up
CP Link [0/0]                up
CP Link [1/2]                down
CP Link [1/1]                down
CP Link [1/0]                down
Crossover LAG                up
CP Link [0/3]                up
CP Link [1/3]                up

```

```

session1@qfabric> show fabric session-host
Identifier: 0281052011000032

```

4. The Director group nonstop software upgrade process continues by downloading and installing software for the fabric manager Routing Engines and the Director devices.

```

root@qfabric>
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing fabric images version 12.2X50-D10.3
Performing cleanup

```

```
Package install complete
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm on peer
Triggering Initial Stage of Fabric Manager Upgrade
Updating CCIF default image to 12.2X50-D10.3
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 15:25:29]: Fabric Manager: Upgrade Initial Stage started
[FM-0 2012-06-05 15:25:38]: FM-0 Master already running on LOCAL DG
[NW-NG-0 2012-06-05 15:25:45]: NW-NG-0 Master already running on LOCAL DG
[FM-0 2012-06-05 15:26:12]: Retrieving package
[FM-0 2012-06-05 15:27:11]: Pushing bundle to re0
[Status 2012-06-05 15:29:06]: Load completed with 0 errors...
[Status 2012-06-05 15:29:06]: Reboot is required to complete upgrade ...
[Status 2012-06-05 15:29:07]: Trying to Connect to Node: FM-0
[Status 2012-06-05 15:29:13]: Rebooting FM-0
[FM-0 2012-06-05 15:29:13]: Waiting for FM-0 to terminate ...
Starting Peer upgrade
```

```
Initiating rolling upgrade of Director peer: version 12.2X50-D10.3
```

```
Inform CCIF regarding rolling upgrade
[Peer Update Status]: Validating install package
jinstall-qfabric-12.2X50-D10.3.rpm
[Peer Update Status]: Cleaning up node for rolling phase one upgrade
[Peer Update Status]: Director group upgrade complete
[Peer Update Status]: COMPLETED
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase
one of rolling upgrade
```

[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade

[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade

[Peer Update Status]: Peer completed phase one of rolling upgrade

- When the system upgrades and reboots the backup Director device DG0, notice how this device is not displayed in the output of the **show fabric administration inventory director-group status** command. Because Director device DG1 appears, this means that the DG1 is operational and acts as the master device.



NOTE: If your second SSH session is being hosted by the rebooting Director device, your session terminates and you need to log back in to establish a new session running on the active Director device.

```
session1@qfabric> show fabric administration inventory director-group status
```

```
Director Group Status Tue Jun 5 15:41:14 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg1	online	master	10.49.215.39	6%	8372272k	4	3 days, 21:25 hrs

Member	Device Id/Alias	Status	Role
dg1	0281052011000032	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:8	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size Used Avail Used% Mounted on

423G	6.0G	395G	2%	/
99M	16M	79M	17%	/boot
93G	7.3G	86G	8%	/pbdata

Director Group Processes

Director Group Manager	online
------------------------	--------


```

Partition Manager          online
Software Mirroring         online
Shared File System master  online
Secure Shell Process       online
Network File System        online
DHCP Server master         online      master

```

```

FTP Server                 online
Syslog                    online
Distributed Management     online
SNMP Trap Forwarder       online
SNMP Process              online
Platform Management       online

```

Interface Link Status

```

-----
Management Interface      up
Control Plane Bridge      up
Control Plane LAG         up
CP Link [0/2]             up
CP Link [0/1]             up
CP Link [0/0]             up
CP Link [1/2]             down
CP Link [1/1]             down
CP Link [1/0]             down
Crossover LAG            up
CP Link [0/3]             up
CP Link [1/3]             up

```

- The upgrade continues with master Director device DG1 suspending CLI services for 15 minutes, transferring mastership to Director device DG0, and then rebooting Director device DG1 (which terminates the CLI session).

root@qfabric>

[Peer Update Status]: Setting peer DG node as the master SFC

Delaying start of local upgrade to allow peer services time to initialize [15 minutes]

Delaying start of local upgrade to allow peer services time to initialize [15 minutes]

Delaying start of local upgrade to allow peer services time to initialize [12 minutes]

Delaying start of local upgrade to allow peer services time to initialize [9 minutes]

Delaying start of local upgrade to allow peer services time to initialize [6 minutes]

Delaying start of local upgrade to allow peer services time to initialize [3 minutes]

[Peer Update Status]: Check for VMs on dg0

Triggering Final Stage of Fabric Manager Upgrade:

Updating FM-0 to Junos version 12.2X50-D10.3

[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started

[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG

[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch

[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3

[DRE-0 2012-06-05 16:12:43]: Retrieving package

[DRE-0 2012-06-05 16:13:46]: ----- re0: -----

[Status 2012-06-05 16:15:17]: Load completed with 0 errors...

[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...

[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...

[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...

```
[DRE-0    2012-06-05 16:18:44]: Running Uptime Test for DRE-0
[DRE-0    2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...
[Status   2012-06-05 16:18:51]: DRE-0 booted successfully ...
Performing post install shutdown and cleanup
```

Broadcast message from root (Tue Jun 5 16:18:51 2012):

The system is going down for reboot NOW!
Director group upgrade complete

```
root@qfabric> Read from remote host qfabric-partition0: Connection reset by
peer
Connection to qfabric-partition0 closed.
```

- Upon reopening the SSH session, notice that Director device DG0 is now the master device hosting the session and Director device DG1 does not appear in the QFabric system inventory while it is rebooting.

```
session1@qfabric> show fabric session-host
Identifier: 0281052011000001
```

```
session1@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun 5 16:21:23 UTC 2012
```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	13%	20739560k	3	36:29 mins

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

Database Server	online
Load Balancer Director	online
QFabric Partition Address	online

Director Group Managed Services

Shared File System	online
Network File System	online
Virtual Machine Server	online
Load Balancer/DHCP	online

Hard Drive Status

Volume ID:4	optimal
Physical ID:1	online
Physical ID:0	online
SCSI ID:1	100%
SCSI ID:0	100%

Size	Used	Avail	Used%	Mounted on
423G	5.3G	396G	2%	/
99M	16M	79M	17%	/boot
93G	7.4G	86G	8%	/pbdata

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager          online
Software Mirroring          online
Shared File System master  online
Secure Shell Process        online
Network File System         online
DHCP Server master         online      master

FTP Server                  online
Syslog                      online
Distributed Management      online
SNMP Trap Forwarder        online
SNMP Process                online
Platform Management        online

Interface Link Status
-----
Management Interface       up
Control Plane Bridge       up
Control Plane LAG          up
CP Link [0/2]              up
CP Link [0/1]              up
CP Link [0/0]              up
CP Link [1/2]              down
CP Link [1/1]              down
CP Link [1/0]              down
Crossover LAG              up
CP Link [0/3]              up
CP Link [1/3]              up

```

8. When Director device DG1 comes back online, it returns to the QFabric system inventory as a backup Director device and hosts some of the Routing Engine processes (which should appear load balanced between the master and backup Director devices).

```

session1@qfabric> show fabric administration inventory director-group status
root@qfabric> show fabric administration inventory director-group status
Director Group Status Tue Jun  5 16:41:02 UTC 2012

```

Member	Status	Role	Mgmt Address	CPU	Free Memory	VMs	Up Time
dg0	online	master	10.49.215.38	15%	14759920k	6	56:09 mins
dg1	online	backup	10.49.215.39	8%	31486680k	0	07:51 mins

Member	Device Id/Alias	Status	Role
dg0	0281052011000001	online	master

Master Services

```

-----
Database Server              online
Load Balancer Director       online
QFabric Partition Address    online

```

Director Group Managed Services

```

-----
Shared File System           online
Network File System          online

```

```

Virtual Machine Server      online
Load Balancer/DHCP         online

Hard Drive Status
-----
Volume ID:4                optimal
Physical ID:1              online
Physical ID:0              online
SCSI ID:1                  100%
SCSI ID:0                  100%

Size  Used Avail Used% Mounted on
----  -
423G  5.3G 396G   2%   /
99M   16M  79M   17%  /boot
93G   7.4G 86G    8%  /pbdata

Director Group Processes
-----
Director Group Manager      online
Partition Manager           online
Software Mirroring          online
Shared File System master   online
Secure Shell Process        online
Network File System         online
DHCP Server master          online      master

FTP Server                   online
Syslog                       online
Distributed Management       online
SNMP Trap Forwarder         online
SNMP Process                 online
Platform Management         online

Interface Link Status
-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]               up
CP Link [0/1]               up
CP Link [0/0]               up
CP Link [1/2]               down
CP Link [1/1]               down
CP Link [1/0]               down
Crossover LAG               up
CP Link [0/3]               up
CP Link [1/3]               up

Member Device Id/Alias  Status  Role
-----
dg1      0281052011000032 online  backup

Director Group Managed Services
-----
Shared File System         online
Network File System        online
Virtual Machine Server     online
Load Balancer/DHCP         online

```

Hard Drive Status

```

-----
Volume ID:8                optimal
Physical ID:1              online
Physical ID:0              online
SCSI ID:1                  100%
SCSI ID:0                  100%

```

Size Used Avail Used% Mounted on

```

-----
423G 5.3G 396G 2% /
99M 16M 79M 17% /boot
93G 7.4G 86G 8% /pbdata

```

Director Group Processes

```

-----
Director Group Manager      online
Partition Manager           online
Software Mirroring           online
Shared File System master    online
Secure Shell Process         online
Network File System          online
DHCP Server master           online    backup

FTP Server                   online
Syslog                       online
Distributed Management        online
SNMP Trap Forwarder          online
SNMP Process                  online
Platform Management          online

```

Interface Link Status

```

-----
Management Interface        up
Control Plane Bridge        up
Control Plane LAG           up
CP Link [0/2]                up
CP Link [0/1]                up
CP Link [0/0]                up
CP Link [1/2]                down
CP Link [1/1]                down
CP Link [1/0]                down
Crossover LAG               up
CP Link [0/3]                up
CP Link [1/3]                up

```

session1@qfabric> show fabric administration inventory infrastructure

dg0:

Routing Engine Type	Hostname	PID	CPU-Use(%)
Fabric control	QFabric_default_FC-1_RE0	27906	2.5
Network Node group	QFabric_default_NW-NG-1_RE1	20421	1.8
Fabric manager	FM-0	4211	1.8
Debug Routing Engine	QFabric_DRE	1575	3.3

dg1:

Routing Engine Type CPU-Use(%)	Hostname	PID	
-----	-----	-----	
Fabric control	QFabric_default_FC-0_RE0	5686	2.3
Network Node group	QFabric_default_NW-NG-0_RE0	5866	1.9
Fabric manager	FM-1	572	1.6

Verifying a Fabric Nonstop Software Upgrade

Purpose During the fabric portion of a nonstop software upgrade, you should expect to see both fabric control Routing Engines upgrade first, followed by the upgrade of each Interconnect device one at a time. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

Action 1. In an SSH session to the QFabric CLI, issue the request for the fabric nonstop software upgrade.

```

root@qfabric> request system software nonstop-upgrade fabric
jinstall-qfabric-12.2X50-D10.3.rpm
[FC-0      2012-06-05 16:48:53]: Retrieving package
[FC-1      2012-06-05 16:48:53]: Retrieving package
[IC-F4912 2012-06-05 16:48:59]: Retrieving package
[FC-0      2012-06-05 16:49:51]: ----- re0: -----
[FC-1      2012-06-05 16:49:52]: ----- re0: -----
[IC-F4912 2012-06-05 16:49:54]: ----- re0: -----
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files
[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in
installation
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active
partition
[Status    2012-06-05 16:58:34]: Load completed with 0 errors...
[Status    2012-06-05 16:58:34]: Reboot is required to complete upgrade ...
[Status    2012-06-05 16:58:34]: Trying to Connect to Node: FC-0
[Status    2012-06-05 16:58:39]: Rebooting FC-0
[Status    2012-06-05 16:58:39]: Trying to Connect to Node: FC-1
[Status    2012-06-05 16:58:44]: Rebooting FC-1
[Status    2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912

```

```
[Status 2012-06-05 16:58:50]: Rebooting IC-F4912
Success
```

- When the fabric components reboot, they appear as **Disconnected** in the output of the **show fabric administration inventory infrastructure fabric-controls** and **show fabric administration inventory interconnect-devices** commands.

```
session1@qfabric> show fabric administration inventory infrastructure fabric-controls
Item                               Identifier                               Connection                               Configuration
Fabric control
  FC-0                               Disconnected
  FC-1                               Disconnected
```

```
session1@qfabric> show fabric administration inventory interconnect-devices IC-F4912
Item                               Identifier                               Connection                               Configuration
Interconnect device
  IC-F4912                               Disconnected
  F4912/RE0                               Disconnected
```

- When the fabric components return to full service, they appear as **Connected** in the output of the **show fabric administration inventory** command.

```
session1@qfabric> show fabric administration inventory
Item                               Identifier                               Connection                               Configuration
Node group
  NW-NG-0                               Connected                               Configured

  P1507-C                               Connected

  RSNG                               Connected                               Configured

  P1550-C                               Connected

  P1571-C                               Connected

Interconnect device
  IC-F4912                               Connected                               Configured

  F4912/RE0                               Connected

Fabric manager
  FM-0                               Connected                               Configured

Fabric control
  FC-0                               Connected                               Configured

  FC-1                               Connected                               Configured

Diagnostic routing engine
  DRE-0                               Connected                               Configured
```

Verifying a Redundant Server Node Group Nonstop Software Upgrade

- Purpose** During the redundant server Node group portion of a nonstop software upgrade, you should expect to see the backup Node device upgrade first, followed by the upgrade of the master Node device. Server Node groups with a single device upgrade the device in the same way as a standalone switch. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.

- Action** 1. In an SSH session to the QFabric CLI, issue the request for the redundant server Node group nonstop software upgrade.

```
root@qfabric> request system software nonstop-upgrade node-group RSNG
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): RSNG
```

```
[RSNG 2012-06-05 17:26:44]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[RSNG 2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active
partition
[RSNG 2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG 2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[RSNG 2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active
partition
```



```

[RSNG 2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG 2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG 2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG 2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG 2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG 2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG 2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG 2012-06-05 17:54:42]: Finished processing all upgrade groups, last
group :1
[RSNG 2012-06-05 17:54:48]: Preparing for Switchover
[RSNG 2012-06-05 17:55:38]: Switchover Completed
[Status 2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success

```

2. Issue the **show system software upgrade status** command to view the status of the upgrade.

```

root@qfabric> show system software upgrade status
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
                RSNG in progress

```

3. During the redundant server Node group upgrade, the backup Node device (in this case, P1571-C) is upgraded first and appears in the **Disconnected** state in the output of the **show fabric administration inventory** command.

```

session1@qfabric> show fabric administration inventory

```

Item	Identifier	Connection	Configuration
Node group			
	NW-NG-0	Connected	Configured
	P1507-C	Connected	
	RSNG	Connected	Configured
	P1550-C	Connected	
	P1571-C	Disconnected	
Interconnect device			
	IC-F4912	Connected	Configured
	F4912/RE0	Connected	
Fabric manager			
	FM-0	Connected	Configured
Fabric control			
	FC-0	Connected	Configured
	FC-1	Connected	Configured
Diagnostic routing engine			
	DRE-0	Connected	Configured

4. After the backup Node device comes back online, the master Node device (in this case, P1550-C) appears in the **Disconnected** state in the output of the **show fabric administration inventory** command while the master Node device upgrades its software.

```

session1@qfabric> show fabric administration inventory

```

Item	Identifier	Connection	Configuration
Node group			

NW-NG-0	Connected	Configured
P1507-C	Connected	
RSNG	Connected	Configured
P1550-C P1571-C	Disconnected Connected	
Interconnect device IC-F4912	Connected	Configured
F4912/RE0	Connected	
Fabric manager FM-0	Connected	Configured
Fabric control FC-0	Connected	Configured
FC-1	Connected	Configured
Diagnostic routing engine DRE-0	Connected	Configured

5. After both Node devices in the redundant server Node group come back online, both Node devices appear as **Connected** to indicate the successful completion of the Node group nonstop software upgrade step.

```
session1@qfabric> show fabric administration inventory
```

Item	Identifier	Connection	Configuration
Node group			
NW-NG-0		Connected	Configured
P1507-C		Connected	
RSNG		Connected	Configured
P1550-C P1571-C		Connected Connected	
Interconnect device IC-F4912		Connected	Configured
F4912/RE0		Connected	
Fabric manager FM-0		Connected	Configured
Fabric control FC-0		Connected	Configured
FC-1		Connected	Configured
Diagnostic routing engine DRE-0		Connected	Configured

Verifying a Network Node Group Nonstop Software Upgrade

Purpose During the network Node group portion of a nonstop software upgrade, you should expect to see the backup network Node group Routing Engine upgrade first, followed by the Node devices within the network Node group upgrading one at a time, and ending with the upgrade of the master network Node group Routing Engine. In addition to the steps below, you can issue the **show system software upgrade status** command to view the progress of the upgrade.



NOTE: If you configure an upgrade group for Node groups containing 2 or more Node devices, all Node devices within the upgrade group reboot at the same time.

Action 1. In an SSH session to the QFabric CLI, issue the request for the network Node group nonstop software upgrade.

```
root@qfabric> request system software nonstop-upgrade node-group NW-NG-0
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): NW-NG-0
```

```
[NW-NG-0 2012-06-01 09:45:06]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
[NW-NG-0 2012-06-01 09:45:06]: Retrieving package
[NW-NG-0 2012-06-01 09:46:18]: Pushing bundle to fpc0
[NW-NG-0 2012-06-01 09:46:52]: fpc0: Validate package...
[NW-NG-0 2012-06-01 09:53:26]: ----- fpc0 -----
[NW-NG-0 2012-06-01 09:54:01]: Step 1 of 20 Creating temporary file system
[NW-NG-0 2012-06-01 09:54:01]: Step 2 of 20 Determining installation source
[NW-NG-0 2012-06-01 09:54:02]: Step 3 of 20 Processing format options
[NW-NG-0 2012-06-01 09:54:02]: Step 4 of 20 Determining installation slice
[NW-NG-0 2012-06-01 09:54:02]: Step 5 of 20 Creating and labeling new slices
[NW-NG-0 2012-06-01 09:54:03]: Step 6 of 20 Create and mount new file system
[NW-NG-0 2012-06-01 09:54:10]: Step 7 of 20 Getting OS bundles
[NW-NG-0 2012-06-01 09:54:10]: Step 8 of 20 Updating recovery media
[NW-NG-0 2012-06-01 09:54:31]: Step 9 of 20 Extracting incoming image
[NW-NG-0 2012-06-01 09:55:43]: Step 10 of 20 Unpacking OS packages
[NW-NG-0 2012-06-01 09:55:46]: Step 11 of 20 Mounting jbase package
[NW-NG-0 2012-06-01 09:56:09]: Step 12 of 20 Creating base OS symbolic links
[NW-NG-0 2012-06-01 09:57:05]: Step 13 of 20 Creating fstab
[NW-NG-0 2012-06-01 09:57:05]: Step 14 of 20 Creating new system files
[NW-NG-0 2012-06-01 09:57:05]: Step 15 of 20 Adding jbundle package
[NW-NG-0 2012-06-01 09:59:30]: Step 16 of 20 Backing up system data
[NW-NG-0 2012-06-01 09:59:44]: Step 17 of 20 Setting up shared partition data
[NW-NG-0 2012-06-01 09:59:44]: Step 18 of 20 Checking package sanity in
installation
[NW-NG-0 2012-06-01 09:59:44]: Step 19 of 20 Unmounting and cleaning up
temporary file systems
[NW-NG-0 2012-06-01 09:59:47]: Step 20 of 20 Setting da0s1 as new active
partition
[NW-NG-0 2012-06-01 09:59:55]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-dc-re.tgz
[NW-NG-0 2012-06-01 09:59:55]: Retrieving package
[NW-NG-0 2012-06-01 10:01:04]: Pushing bundle to re1
[NW-NG-0 2012-06-01 10:01:35]: re1: Validate package...
```

```
[NW-NG-0 2012-06-01 10:02:56]: re0: Validate package...
[NW-NG-0 2012-06-01 10:04:45]: Rebooting Backup RE
[NW-NG-0 2012-06-01 10:08:31]: Initiating Chassis In-Service-Upgrade
[NW-NG-0 2012-06-01 10:08:52]: Upgrading group: 0 fpc: 0
[NW-NG-0 2012-06-01 10:18:33]: Upgrade complete for group:0
[NW-NG-0 2012-06-01 10:18:33]: Finished processing all upgrade groups, last
group :0
[NW-NG-0 2012-06-01 10:18:37]: Preparing for Switchover
[NW-NG-0 2012-06-01 10:18:55]: Switchover Completed
[Status 2012-06-01 10:18:58]: Upgrade completed with 0 errors
Success
```

2. Issue the **show system software upgrade status** command to view the status of the upgrade.

```
root@qfabric> show system software upgrade status
Wed Jan 16 22:06:02 2013 Software nonstop upgrade on:
NW-NG-0 in progress
```

3. Verify the progress of the upgrade by issuing the **show chassis nonstop-upgrade node-group**, **show fabric administration inventory**, **show fabric administration inventory infrastructure**, and **show fabric administration inventory node-groups NW-NG-0** commands. You should see the backup network Node group Routing Engine reboot first, followed by each Node device within the network Node group, and ending with the reboot of master network Node group Routing Engine. Restarting devices appear as **Disconnected** in the output of the **show fabric administration inventory** command and restarting Routing Engines do not appear in output of the **show fabric administration inventory infrastructure** command until they return to service.

Related Documentation

- [Performing a Nonstop Software Upgrade on the QFabric System on page 36](#)
- [Understanding Nonstop Software Upgrade for QFabric Systems on page 31](#)
- [show chassis nonstop-upgrade node-group on page 170](#)
- *show fabric administration inventory*
- *show fabric administration inventory director-group status*
- *show fabric administration inventory infrastructure*
- *show fabric administration inventory interconnect-devices*
- *show fabric administration inventory node-groups*

PART 4

Virtual Router Redundancy Protocol

- [Using VRRP on page 63](#)

CHAPTER 4

Using VRRP

- [Understanding VRRP on page 63](#)
- [Understanding VRRP Between QFabric Systems on page 66](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)
- [Configuring Basic VRRP Support on page 74](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 75](#)
- [Configuring the Startup Period for VRRP Operations on page 76](#)
- [Configuring the Advertisement Interval for the VRRP Master on page 76](#)
- [Configuring VRRP Preemption and Hold Time on page 77](#)
- [Configuring a Route to Be Tracked on page 79](#)
- [Configuring a Logical Interface to Be Tracked on page 79](#)
- [Configuring a Backup to Accept Packets Destined for the Virtual IP Address on page 81](#)
- [Configuring Passive ARP Learning for VRRP Backups on page 82](#)
- [Configuring the Silent Period on page 82](#)
- [Configuring Inheritance for a VRRP Group on page 83](#)

Understanding VRRP

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

- [Overview of VRRP on page 63](#)
- [Sample VRRP Topology on page 64](#)

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 3 on page 65](#) 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 3: Basic VRRP Topology



The following illustrates basic VRRP behavior using [Figure 3 on page 65](#) 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 on page 74](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)
- [Understanding VRRP Between QFabric Systems on page 66](#)

Understanding VRRP Between QFabric Systems

Juniper Networks QFabric systems support the Virtual Router Redundancy Protocol (VRRP). This topic covers:

- [VRRP Differences on QFabric Systems on page 66](#)
- [Configuration Details on page 66](#)

VRRP Differences on QFabric Systems

Configuring servers on your network with static routes to a default gateway minimizes configuration effort and complexity and reduces processing overhead. However, a failure of the default gateway normally results in a catastrophic event, isolating the servers. Using Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for servers if the primary gateway fails.

Switches configured with VRRP share a virtual IP (VIP) address, which is the address you configure as the default route on the servers. In normal VRRP operation, one of the switches is the VRRP master, meaning that it owns the VIP 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. If the master fails, the backup switch with the highest priority becomes the master and takes ownership of the VIP within a few seconds. This is done without any interaction with the servers.

You can configure two QFabric systems to participate in a VRRP configuration as if they were two standalone switches. However, in normal VRRP operation, only one system can be the master for a given VRRP group at any one time, which means that only one system can act as a default gateway using the VIP configured for the group. When running VRRP over two QFabric systems, you might want both systems to simultaneously use the VIP to act as a gateway and forward traffic. To achieve this, you can configure a firewall filter to block the VRRP advertisement packets between the QFabric systems on the link between the two network Node groups. When you do this, both QFabric systems act as master and forward traffic received by the VIP (which is the default gateway address that you configure on servers connected to both QFabric systems). If you use VMware's vMotion, this configuration allows virtual machines to transition between servers connected to the QFabric systems without updating their default gateway information. For example, a virtual machine running on a server connected to a QFabric system in data center A can transition to a server connected to a QFabric system in data center B without needing to resolve a new gateway IP address and MAC address because both QFabric systems use the same VIP.



NOTE: To use a firewall filter to block VRRP traffic, create a firewall term that matches traffic for protocol vrrp and discards that traffic.

Configuration Details

Configuring a VRRP group across two QFabric systems is similar to configuring VRRP on two switches. The main differences are listed here:

- All the interfaces in both QFabric systems that participate in VRRP must be members of the same VLAN.
- You must create routed VLAN interfaces (RVIs) in that VLAN on both QFabric systems.
- The IP addresses that you assign to both RVIs must be in the same subnet.
- You must configure VRRP on the RVIs.
- Both RVIs must be members of the same VRRP group. This is what allows the two QFabric systems to share a virtual IP address.

The following tables list the elements of an example VRRP configuration running on two QFabric systems—QFabric system A and QFabric system B. This example is configured so that both QFabric systems act as the VRRP master for VIP 10.1.1.50/24 and assumes that a firewall filter blocks the VRRP advertisements between the systems.

[Table 5 on page 67](#) lists the required characteristics of the RVIs in the example configuration.



NOTE: Most of the configuration settings in the following tables would also apply in a traditional VRRP configuration. However, the advertisement interval and priority settings would need to be different (as noted).

Table 5: RVIs on QFabric systems in example VRRP configuration

RVI on QFabric System A	RVI on QFabric System B
vlan.100	vlan.200
Member of VLAN 100. (Note that the VLAN is the same on both QFabric systems.)	Member of VLAN 100
IP address 10.1.1.100/24	IP address 10.1.1.200/24
Member of VRRP group 500	Member of VRRP group 500
Virtual IP address 10.1.1.50/24	Virtual IP address 10.1.1.50/24

You must configure VRRP on the RVIs on both QFabric systems. [Table 6 on page 67](#) lists the elements of a sample VRRP configuration on each RVI. Note that with the exception of the priority, the parameters *must* be the same on both systems.

Table 6: Sample VRRP configuration each RVI

VRRP on RVI on QFabric System A	VRRP on RVI on QFabric System B
VRRP group 500	VRRP group 500
Virtual IP address 10.1.1.50/24	Virtual IP address 10.1.1.50/24

Table 6: Sample VRRP configuration each RVI (*continued*)

Advertisement interval 60 seconds. (In a normal VRRP configuration, you would set this interval to be much smaller, such as 1 second. However, in this configuration these packets are blocked by the firewall filter on the interface that connects to QFabric system B, so there is no need to send them frequently.)	Advertisement interval 60 seconds
Authentication type md5	Authentication type md5
Authentication key \$9\$1.4ElMVb2aGi4aZjkqzFRhSeWx7-wY2aM8	Authentication key \$9\$1.4ElMVb2aGi4aZjkqzFRhSeWx7-wY2aM8
Priority 254. (In a normal VRRP configuration, this value would be different on the two systems and the system with the higher value would be the master. However, in this configuration both systems are acting as master, so you do not have to configure different values.)	Priority 254



NOTE: Priority 255 is not supported for RVIs.

[Table 7 on page 68](#) lists the all the interfaces on QFabric system A in the example configuration and identifies what they connect to.

Table 7: Interfaces on QFabric system A. All interfaces are members of VLAN 100.

VLAN 100 Interfaces on QFabric System A	Connects To
vlan.100	vlan.200
Network Node group interface QFA-NNG:xe-0/0/0	QFB-NNG:xe-0/0/0 on QFabric system B
Network Node group interface QFA-NNG:xe-0/0/1	Redundant server Node group interface QFA-RSNG:xe-0/0/0
Redundant server Node group interface QFA-RSNG:xe-0/0/0	Connects to a network Node group interface QFA-NNG:xe-0/0/1
Redundant server Node group interface QFA-RSNG:xe-0/0/1	LAN with servers running virtual machines

[Table 8 on page 68](#) lists the all the interfaces on QFabric system B in the example configuration and identifies what they connect to.

Table 8: Interfaces on QFabric system B. All interfaces are members of VLAN 100 (same as on QFabric system A).

VLAN 100 Interfaces on QFabric System B	Connects To
vlan.200	vlan.100
Network Node group interface QFB-NNG:xe-0/0/0	QFA-NNG:xe-0/0/0 on QFabric system A
Network Node group interface QFB-NNG:xe-0/0/1	Redundant server Node group interface QFB-RSNG:xe-0/0/0

Table 8: Interfaces on QFabric system B. All interfaces are members of VLAN 100 (same as on QFabric system A). (*continued*)

Redundant server Node group interface QFB-RSNG:xe-0/0/0	Connects to a network Node group interface QFB-NNG:xe-0/0/1
Redundant server Node group interface QFB-RSNG:xe-0/0/1	LAN with servers running virtual machines

- Related Documentation**
- [Understanding VRRP on page 63](#)
 - [Configuring Basic VRRP Support on page 74](#)
 - [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 69](#)
- [Overview and Topology on page 69](#)
- [Configuring VRRP on Both Switches on page 70](#)
- [Verification on page 73](#)

Requirements

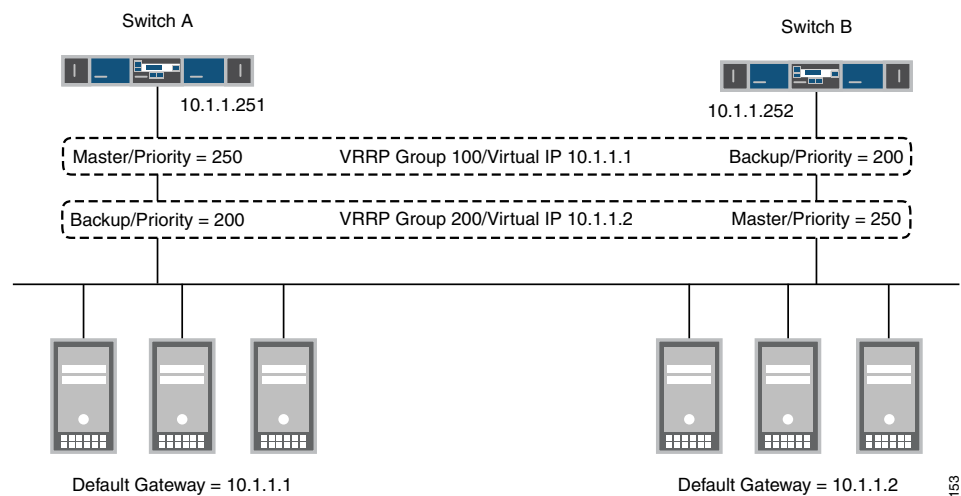
This example uses the following hardware and software components:

- 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 4 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 4: VRRP Load-Sharing Configuration



9041153

This example shows a simple configuration to illustrate the basic steps for configuring two switches running VRRP to back each other up. [Table 9 on page 70](#) lists VRRP settings for each switch.

Table 9: Settings for VRRP Load-Sharing Example

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

1. 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
```

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 10.1.1.251/24 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 73](#)

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 **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 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 63](#)
 - [Configuring Basic VRRP Support on page 74](#)

Configuring Basic VRRP Support

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 63](#)
- [Configuring the Startup Period for VRRP Operations on page 76](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 75](#)

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 63](#)
- [Configuring Basic VRRP Support on page 74](#)

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 63](#)
- [Configuring Basic VRRP Support on page 74](#)
- [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 77](#)
- [Modifying the Advertisement Interval in Milliseconds on page 77](#)

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 63](#)
- [Configuring Basic VRRP Support on page 74](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

Configuring VRRP Preemption and Hold Time

- [Configuring VRRP Preemption on page 78](#)
- [Configuring the Preemption Hold Time on page 78](#)
- [Overriding the Hold Time on page 78](#)

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 63](#)

- [Configuring Basic VRRP Support on page 74](#)
- [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 63](#)
- [Configuring Basic VRRP Support on page 74](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)
- [Configuring a Logical Interface to Be Tracked on page 79](#)

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 10 on page 81](#):

Table 10: 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 63](#)
- [Configuring Basic VRRP Support on page 74](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)
- [Configuring a Route to Be Tracked on page 79](#)

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 63](#)
 - [Configuring Basic VRRP Support on page 74](#)
 - [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 63](#)
 - [Configuring Basic VRRP Support on page 74](#)
 - [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 63](#)
 - [Configuring Basic VRRP Support on page 74](#)

- [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 63](#)
- [Configuring Basic VRRP Support on page 74](#)
- [Example: Configuring VRRP for Load Sharing on page 69](#)

PART 5

Configuration Statements and Operational Commands

- [Configuration Statements for Graceful Restart on page 87](#)
- [Configuration Statement for Graceful Switchover on page 103](#)
- [Configuration Statements for VRRP on page 105](#)
- [Operational Commands for Graceful Restart on page 129](#)
- [Operational Command for Graceful Switchover on page 155](#)
- [Operational Commands for Nonstop Software Upgrade on page 163](#)
- [Operational Commands for VRRP on page 171](#)

CHAPTER 5

Configuration Statements for Graceful Restart

- `disable` on page 88
- `disable` (BGP Graceful Restart) on page 89
- `graceful-restart` (Enabling Globally) on page 90
- `graceful-restart` (Protocols BGP) on page 92
- `graceful-restart` (Protocols OSPF) on page 93
- `helper-disable` (OSPF) on page 95
- `no-strict-lsa-checking` on page 96
- `notify-duration` on page 97
- `redundancy` (Graceful Switchover) on page 98
- `restart-duration` on page 99
- `restart-time` (BGP Graceful Restart) on page 100
- `stale-routes-time` on page 101

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"> • Enabling Graceful Restart • Configuring Routing Protocols Graceful Restart on page 8 • 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 on page 4

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	routing—To view this statement in the configuration.
Level	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 8 • Configuring Graceful Restart for QFabric Systems on page 4 • graceful-restart on page 92
------------------------------	---

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	[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]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced for QFX switches in Junos OS Release 13.2X51-D15 Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	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.




NOTE:

- For VPNs, the **graceful-restart** 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 **graceful-restart** configurations change.
-

Default	Graceful restart is disabled by default.
Options	The remaining statements are explained separately.
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 Routing Protocols Graceful Restart on page 8](#)
 - *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 on page 4](#)

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.</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>  <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> <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 8 • Configuring Graceful Restart for QFabric Systems on page 4 • <i>Junos OS High Availability Library for Routing Devices</i>

graceful-restart (Protocols OSPF)

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> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>
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**
- *Example: Configuring Graceful Restart for OSPF*
 - *Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart*
 - *Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart*
 - *Example: Disabling Strict LSA Checking for OSPF Graceful Restart*
 - *Junos OS High Availability Library for Routing Devices*

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 8 • 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 8• Configuring Graceful Restart for QFabric Systems on page 4• <i>maximum-neighbor-recovery-time</i>• <i>recovery-time</i>

notify-duration

Syntax	<code>notify-duration <i>seconds</i>;</code>
Hierarchy Level	<p>[edit protocols (ospf ospf3) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart],</p> <p>[edit routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</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	Specify the length of time the router or switch notifies helper OSPF routers that it has completed graceful restart.
Options	<p><i>seconds</i>—Length of time in the router notifies helper OSPF routers that it has completed graceful restart.</p> <p>Range: 1 through 3600</p> <p>Default: 30</p>
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"> • Configuring Graceful Restart Options for OSPF and OSPFv3 on page 8 • Configuring Graceful Restart for QFabric Systems on page 4 • restart-duration on page 99

redundancy (Graceful Switchover)

Syntax	<pre>redundancy { failover { on-disk-failure; on-loss-of-keepalives; } graceful-switchover; }</pre>
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• graceful-switchover on page 103• <i>Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)</i>• Configuring Graceful Routing Engine Switchover on page 26• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>• <i>High Availability Features for EX Series Switches Overview</i>

restart-duration

Syntax	<code>restart-duration <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (isis ospf ospf3 pim) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</p> <p>[edit protocols (esis isis ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) 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 14.1X53-D20 for the OCX Series.</p>
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><i>seconds</i>—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 Routing Protocols Graceful Restart on page 8](#)
- [Configuring Graceful Restart for MPLS-Related Protocols](#)
- [Configuring VPN Graceful Restart](#)
- [Configuring Graceful Restart for VPNs](#)
- [Configuring Logical System Graceful Restart](#)
- [Graceful Restart Configuration Statements](#)

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 8](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 12](#)
- [Configuring Graceful Restart for QFabric Systems on page 4](#)
- [stale-routes-time on page 101](#)

stale-routes-time

Syntax	<code>stale-routes-time</code> <i>seconds</i> ;
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 8 • Configuring Graceful Restart for QFabric Systems on page 4 • restart-time (BGP Graceful Restart) on page 100

CHAPTER 6

Configuration Statement for Graceful Switchover

- [graceful-switchover](#) on page 103
- [redundancy \(Graceful Switchover\)](#) on page 104

[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">• <i>Example: Configuring Nonstop Active Routing on Switches</i>• Configuring Graceful Routing Engine Switchover on page 26• <i>Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)</i>• <i>Configuring Nonstop Active Routing on Switches</i>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>

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 103• <i>Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)</i>• Configuring Graceful Routing Engine Switchover on page 26• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>• <i>High Availability Features for EX Series Switches Overview</i>

CHAPTER 7


Configuration Statements for VRRP

- [accept-data on page 106](#)
- [advertise-interval on page 107](#)
- [asymmetric-hold-time on page 108](#)
- [authentication-key on page 109](#)
- [authentication-type on page 110](#)
- [bandwidth-threshold on page 111](#)
- [failover-delay on page 112](#)
- [fast-interval on page 113](#)
- [hold-time \(VRRP\) on page 114](#)
- [interface \(VRRP Group\) on page 115](#)
- [preempt \(VRRP\) on page 116](#)
- [priority \(Protocols VRRP\) on page 117](#)
- [priority-cost \(VRRP\) on page 118](#)
- [priority-hold-time on page 119](#)
- [route \(Interfaces\) on page 120](#)
- [startup-silent-period on page 121](#)
- [traceoptions on page 122](#)
- [track \(VRRP\) on page 124](#)
- [virtual-address on page 125](#)
- [vrrp-group on page 126](#)

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> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a router that is acting as the master router accepts all packets destined for the virtual IP address.</p> <ul style="list-style-type: none"> • accept-data—Enable the master router to accept all packets destined for the virtual IP address. • no-accept-data—Prevent the master router from accepting packets other than the ARP packets destined for the virtual IP address.
Default	<p>If the router acting as the master router is the IP address owner or has its priority set to 255, the master router, by default, responds to all packets sent to the virtual IP address. However, if the router acting as the master router does not own the IP address or has its priority set to a value less than 255, the master router responds only to ARP requests.</p>
<div>  <div> <p>NOTE:</p> <ul style="list-style-type: none"> • If you want to restrict the incoming IP packets to ICMP packets 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>). </div> </div>	
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 113 • <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 77

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> 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 inet address <i>address</i> vrrp-group <i>group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX 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) on page 75 • authentication-type on page 110 • version-3

authentication-type

Syntax	<code>authentication-type <i>authentication</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	<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> NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups.</div>	
Options	<p>authentication—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	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 75authentication-key on page 109version-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> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX 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 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.</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 79

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">• <i>Troubleshooting VRRP</i>• show vrrp on page 172

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 Release 14.1x53-D20 for the OCX 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 76 • advertise-interval on page 107 • advertise-interval on page 107 • <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. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX 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	seconds —Hold-time period. Range: 0 through 3600 seconds Default: 0 seconds (VRRP preemption is not timed.)
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —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 77

interface (VRRP Group)


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> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.



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.

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 explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked for a VRRP Group • Configuring a Logical Interface to Be Tracked on page 79 • Junos OS Services Interfaces Library for Routing Devices

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>  <p>NOTE: By default, a higher-priority backup router can preempt a lower-priority master router.</p> <p>.....</p> </div> <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 77


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 on page 74

priority-cost (VRRP)

Syntax	<code>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</i> track interface <i>interface-name</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</i> track interface <i>interface-name</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</i> track interface <i>interface-name</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</i> track interface <i>interface-name</i>]</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	<i>priority</i> —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 79

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 style="display: flex; align-items: center;">  <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> </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 79

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 79

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	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 the Startup Period for VRRP Operations</i> • Configuring the Startup Period for VRRP Operations on page 76

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 79• Configuring a Route to Be Tracked on page 79

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 on page 74

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 logical-system-name 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:00:01:00 through 00:00:5e:00:01: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 on page 74](#)
 - [Example: Configuring VRRP for Load Sharing on page 69](#)
 - *vrrp-inet6-group*

CHAPTER 8

Operational Commands for Graceful Restart

- [Verifying Graceful Restart Operation on page 129](#)
- [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 129](#)
- [Verifying BGP Graceful Restart on page 130](#)
- [Verifying IS-IS and OSPF Graceful Restart on page 130](#)
- [Verifying CCC and TCC Graceful Restart on page 131](#)

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 192.255.10.1
Peer: 192.255.10.1+179 AS 64595 Local: 192.255.5.1+1106 AS 64595
  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: 192.255.5.1 Holdtime: 90 Preference: 170
  IPsec SA Name: hope
  Number of flaps: 0
  Peer ID: 192.255.10.1    Local ID: 192.255.5.1    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 8](#)).

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 192.255.5.1
Oct 8 05:20:14 Received multiple grace lsa from 192.255.5.1
```

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 on page 4](#)

show bgp neighbor

List of Syntax	Syntax on page 132 Syntax (EX Series Switch, QFX Series, and OCX Series) on page 132
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 139](#)
[show bgp neighbor \(CLNS\) \(Not supported on the OCX Series.\) on page 140](#)
[show bgp neighbor \(Layer 2 VPN\) \(Not supported on the OCX Series.\) on page 141](#)
[show bgp neighbor \(Layer 3 VPN\) \(Not supported on the OCX Series.\) on page 143](#)
[show bgp neighbor neighbor-address on page 143](#)
[show bgp neighbor neighbor-address on page 144](#)
[show bgp neighbor orf neighbor-address detail on page 145](#)

Output Fields [Table 11 on page 133](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 11: 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.

Table 11: 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. • 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 11: 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. • 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. • 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.
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.

Table 11: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
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.
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 for restart configured on peer	Names of address families configured for restart.

Table 11: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
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-as1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
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 .

Table 11: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
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.
Output queue	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. It also specifies the routing table name and the NLRI they represent in the format (routing table name, NLRI). NOTE: The output queues of routing tables that are not advertised, will only show up at extensive output level.
Trace options	Configured tracing of BGP protocol packets and operations.

Table 11: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates rcv	(orf option only) Number of outbound-route filters received for each configured address family. 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.
Immediate	(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. 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.
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: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  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) (Not supported on the OCX Series.)

```

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) (Not supported on the OCX Series.)

```
user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
  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 65100 Local: 10.69.104.1      AS 65104
  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 l2vpn
Table bgp.l3vpn.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.l2vpn.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.l2vpn.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: 4.4.4.4+179 AS 10045 Local: 5.5.5.5+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: 5.5.5.5 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/bgpr 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 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 146 Syntax (QFX Series and OCX Series) on page 146 Syntax (TX Matrix Routers) on page 146
Syntax	<code>show log</code> <code><filename user <username>></code>
Syntax (QFX Series and OCX Series)	<code>show log filename</code> <code><device-type (device-id device-alias)></code>
Syntax (TX Matrix Routers)	<code>show log</code> <code><all-lcc lcc number scc></code> <code><filename user <username>></code>
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 11.1 for the QFX Series. Option <i>device-type (device-id device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.
Options	none —List all log files. <all-lcc lcc number scc> —(TX Matrix routers 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). device-type —(QFabric system only) (Optional) Display log messages for only one of the following device types: <ul style="list-style-type: none">• 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.



NOTE: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

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

List of Sample Output [show log on page 147](#)
[show log filename on page 147](#)
[show log filename \(QFabric System\) on page 148](#)
[show log user on page 148](#)

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 rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv 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_REO_ 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_REO_ 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
darius  mg2546                Thu Oct  1 19:37   still logged in
darius  mg2529                Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518                Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575                Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2   Wed Sep 30 01:03 - 01:22 (00:19)

```

show (ospf | ospf3) overview

List of Syntax	Syntax on page 149 Syntax (EX Series Switch and QFX Series) on page 149
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> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX 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 151 show ospf overview (With Database Protection) on page 152 show ospf3 overview (With Database Protection) on page 152 show ospf overview extensive on page 152
Output Fields	<p>Table 12 on page 150 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear.</p>

Table 12: 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 12: 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 . NOTE: The Authentication Type field refers to the authentication configured at the <code>[edit protocols ospf area area-id]</code> level. Any authentication configured for an interface in this area will not affect the value of this field.	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 9

Operational Command for Graceful Switchover

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

show system switchover

List of Syntax	Syntax on page 156 Syntax (TX Matrix Router) on page 156 Syntax (TX Matrix Plus Router) on page 156 Syntax (MX Series Router) on page 156
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	<p>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.</p> <p>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.</p> <p>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 an error message that graceful Routing Engine switchover (GRES) is not enabled on this member.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Routing Matrix with a TX Matrix Plus Router Solutions Page
List of Sample Output	<p>show system switchover (Backup Routing Engine) on page 159</p> <p>show system switchover all-lcc (Routing Matrix) on page 159</p>
Output Fields	<p>Table 13 on page 158 describes the output fields for the show system switchover command. Output fields are listed in the approximate order in which they appear.</p>

Table 13: show system switchover Output Fields

Field Name	Field Description
Graceful switchover	<p>Display graceful Routing Engine switchover status:</p> <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	<p>State of the configuration database:</p> <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.
Kernel database	<p>State of the kernel database:</p> <ul style="list-style-type: none"> • Ready—Kernel database has synchronized. • 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> <ul style="list-style-type: none"> • Steady State—Peer completed switchover transition. • Peer Connected—Peer in switchover transition.

Sample Output

show system switchover (Backup Routing Engine)

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

show system switchover all-lcc (Routing Matrix)

```
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 160 show task replication (Issued on the Backup Routing Engine) on page 161
Output Fields	Table 14 on page 160 lists the output fields for the show task replication command. Output fields are listed in the approximate order in which they appear.

Table 14: 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 10

Operational Commands for Nonstop Software Upgrade

- `request system software nonstop-upgrade`
- `show chassis nonstop-upgrade node-group`

request system software nonstop-upgrade

Syntax	<code>request system software nonstop-upgrade <i>package-name</i></code> <code><fabric ></code> <code><director-group></code> <code><node-group <i>name</i>></code>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. You should upgrade the devices in the following order: Director group, fabric controls and Interconnect devices, and network and server Node groups.
Options	<p><i>package-name</i>—Location from which the software is to be installed. For example:</p> <ul style="list-style-type: none">• <i>protocol://hostname/pathname/package-name</i>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <i>protocol</i> with one of the following:<ul style="list-style-type: none">• ftp—File Transfer Protocol. Use <i>ftp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>ftp://<username>:<password>@hostname/pathname/package-name</i>. 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 <i>http://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>http://<username>:<password>@hostname/pathname/package-name</i>. 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 <i>scp://hostname/pathname/package-name</i>. To specify authentication credentials, use <i>scp://<username>:<password>@hostname/pathname/package-name</i>.

**NOTE:**

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

director-group—Install software package on the Director group and Fabric managers.

fabric—Install software package on the Interconnect devices and Fabric controls.

node-group *name* —Install software package on the redundant server Node group, server Node group, or network Node group.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • Performing a Nonstop Software Upgrade on the QFabric System on page 36 • Verifying Nonstop Software Upgrade for QFabric Systems on page 41 • show chassis nonstop-upgrade node-group on page 170
List of Sample Output	request system software nonstop-upgrade director-group on page 165 request system software nonstop-upgrade fabric on page 167 request system software nonstop-upgrade node-group (Redundant Server Node Group) on page 167 request system software nonstop-upgrade node-group (Server Node Group) on page 169
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software nonstop-upgrade director-group

```

user@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D10.3.rpm
Validating update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm
Installing fabric images version 12.2X50-D10.3
Performing cleanup
Package install complete
Installing update package jinstall-qfabric-12.2X50-D10.3.rpm on peer
Triggering Initial Stage of Fabric Manager Upgrade
Updating CCIF default image to 12.2X50-D10.3
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 15:25:29]: Fabric Manager: Upgrade Initial Stage started
[FM-0 2012-06-05 15:25:38]: FM-0 Master already running on LOCAL DG
[NW-NG-0 2012-06-05 15:25:45]: NW-NG-0 Master already running on LOCAL DG
[FM-0 2012-06-05 15:26:12]: Retrieving package
[FM-0 2012-06-05 15:27:11]: Pushing bundle to re0
[Status 2012-06-05 15:29:06]: Load completed with 0 errors...
[Status 2012-06-05 15:29:06]: Reboot is required to complete upgrade ...
[Status 2012-06-05 15:29:07]: Trying to Connect to Node: FM-0
[Status 2012-06-05 15:29:13]: Rebooting FM-0
[FM-0 2012-06-05 15:29:13]: Waiting for FM-0 to terminate ...
Starting Peer upgrade

Initiating rolling upgrade of Director peer: version 12.2X50-D10.3

Inform CCIF regarding rolling upgrade
[Peer Update Status]: Validating install package jinstall-qfabric-12.2X50-D10.3.rpm
[Peer Update Status]: Cleaning up node for rolling phase one upgrade
[Peer Update Status]: Director group upgrade complete
[Peer Update Status]: COMPLETED
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling
upgrade

```

```
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to return after reboot and start phase one of rolling upgrade
[Peer Update Status]: Waiting for peer to complete phase one of rolling upgrade
[Peer Update Status]: Peer completed phase one of rolling upgrade
Setting peer DG node as the master SFC
```

```
Delaying start of local upgrade to allow peer services time to initialize [15 minutes]
Delaying start of local upgrade to allow peer services time to initialize [15 minutes]
Delaying start of local upgrade to allow peer services time to initialize [12 minutes]
Delaying start of local upgrade to allow peer services time to initialize [9 minutes]
Delaying start of local upgrade to allow peer services time to initialize [6 minutes]
Delaying start of local upgrade to allow peer services time to initialize [3 minutes]
[Peer Update Status]: Check for VMs on dg0
Triggering Final Stage of Fabric Manager Upgrade:
Updating FM-0 to Junos version 12.2X50-D10.3
[Status 2012-06-05 16:10:12]: Fabric Manager: Upgrade Final Stage started
[NW-NG-0 2012-06-05 16:10:22]: Transferring NW-NG-0 Mastership to REMOTE DG
[NW-NG-0 2012-06-05 16:11:44]: Finished NW-NG-0 Mastership switch
[Status 2012-06-05 16:11:45]: Upgrading FM-0 VM on worker DG to 12.2X50-D10.3
[DRE-0 2012-06-05 16:12:43]: Retrieving package
[DRE-0 2012-06-05 16:13:46]: ----- re0: -----
[Status 2012-06-05 16:15:17]: Load completed with 0 errors...
[Status 2012-06-05 16:15:17]: Reboot is required to complete upgrade ...
[DRE-0 2012-06-05 16:15:22]: Waiting for DRE-0 to terminate ...
[DRE-0 2012-06-05 16:15:34]: Waiting for DRE-0 to come back ...
[DRE-0 2012-06-05 16:18:44]: Running Uptime Test for DRE-0
[DRE-0 2012-06-05 16:18:51]: Uptime Test for DRE-0 Passed ...
```

```
[Status 2012-06-05 16:18:51]: DRE-0 booted successfully ...
Performing post install shutdown and cleanup
```

```
Broadcast message from root (Tue Jun 5 16:18:51 2012):
```

```
The system is going down for reboot NOW!
Director group upgrade complete
```

```
root@qfabric> Read from remote host qfabric-partition0: Connection reset by peer
Connection to qfabric-partition0 closed.
```

request system software nonstop-upgrade fabric

```
user@qfabric> request system software nonstop-upgrade fabric
jinstall-qfabric-12.2X50-D10.3.rpm
[FC-0 2012-06-05 16:48:53]: Retrieving package
[FC-1 2012-06-05 16:48:53]: Retrieving package
[IC-F4912 2012-06-05 16:48:59]: Retrieving package
[FC-0 2012-06-05 16:49:51]: ----- re0: -----
[FC-1 2012-06-05 16:49:52]: ----- re0: -----
[IC-F4912 2012-06-05 16:49:54]: ----- re0: -----
[IC-F4912 2012-06-05 16:50:42]: Step 1 of 20 Creating temporary file system
[IC-F4912 2012-06-05 16:50:42]: Step 2 of 20 Determining installation source
[IC-F4912 2012-06-05 16:50:43]: Step 3 of 20 Processing format options
[IC-F4912 2012-06-05 16:50:43]: Step 4 of 20 Determining installation slice
[IC-F4912 2012-06-05 16:50:43]: Step 5 of 20 Creating and labeling new slices
[IC-F4912 2012-06-05 16:50:44]: Step 6 of 20 Create and mount new file system
[IC-F4912 2012-06-05 16:50:53]: Step 7 of 20 Getting OS bundles
[IC-F4912 2012-06-05 16:50:53]: Step 8 of 20 Updating recovery media
[IC-F4912 2012-06-05 16:51:17]: Step 9 of 20 Extracting incoming image
[IC-F4912 2012-06-05 16:52:56]: Step 10 of 20 Unpacking OS packages
[IC-F4912 2012-06-05 16:52:59]: Step 11 of 20 Mounting jbase package
[IC-F4912 2012-06-05 16:53:28]: Step 12 of 20 Creating base OS symbolic links
[IC-F4912 2012-06-05 16:54:45]: Step 13 of 20 Creating fstab
[IC-F4912 2012-06-05 16:54:45]: Step 14 of 20 Creating new system files
[IC-F4912 2012-06-05 16:54:46]: Step 15 of 20 Adding jbundle package
[IC-F4912 2012-06-05 16:58:15]: Step 16 of 20 Backing up system data
[IC-F4912 2012-06-05 16:58:18]: Step 17 of 20 Setting up shared partition data
[IC-F4912 2012-06-05 16:58:18]: Step 18 of 20 Checking package sanity in
installation
[IC-F4912 2012-06-05 16:58:18]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[IC-F4912 2012-06-05 16:58:22]: Step 20 of 20 Setting da0s1 as new active partition
[Status 2012-06-05 16:58:34]: Load completed with 0 errors...
[Status 2012-06-05 16:58:34]: Reboot is required to complete upgrade ...
[Status 2012-06-05 16:58:34]: Trying to Connect to Node: FC-0
[Status 2012-06-05 16:58:39]: Rebooting FC-0
[Status 2012-06-05 16:58:39]: Trying to Connect to Node: FC-1
[Status 2012-06-05 16:58:44]: Rebooting FC-1
[Status 2012-06-05 16:58:44]: Trying to Connect to Node: IC-F4912
[Status 2012-06-05 16:58:50]: Rebooting IC-F4912
Success
```

request system software nonstop-upgrade node-group (Redundant Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group RSNG
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): RSNG

[RSNG 2012-06-05 17:26:44]: Starting with package
ftp://169.254.0.3/pub/images/12.2X50-D10.3/jinstall-qfx.tgz
```

```
[RSNG 2012-06-05 17:26:44]: Retrieving package
[RSNG 2012-06-05 17:28:56]: Pushing bundle to fpc1
[RSNG 2012-06-05 17:29:26]: fpc1: Validate package...
[RSNG 2012-06-05 17:35:22]: fpc0: Validate package...
[RSNG 2012-06-05 17:35:49]: ----- fpc1 -----
[RSNG 2012-06-05 17:36:25]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:36:26]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:36:26]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:36:26]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:36:27]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:36:27]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:36:35]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:36:35]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:36:56]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:38:07]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:38:16]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:38:41]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:39:41]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:39:42]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:39:42]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:16]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:32]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:33]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:33]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG 2012-06-05 17:42:36]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG 2012-06-05 17:42:51]: ----- fpc0 - master -----
[RSNG 2012-06-05 17:42:51]: Step 1 of 20 Creating temporary file system
[RSNG 2012-06-05 17:42:51]: Step 2 of 20 Determining installation source
[RSNG 2012-06-05 17:42:51]: Step 3 of 20 Processing format options
[RSNG 2012-06-05 17:42:51]: Step 4 of 20 Determining installation slice
[RSNG 2012-06-05 17:42:51]: Step 5 of 20 Creating and labeling new slices
[RSNG 2012-06-05 17:42:51]: Step 6 of 20 Create and mount new file system
[RSNG 2012-06-05 17:42:51]: Step 7 of 20 Getting OS bundles
[RSNG 2012-06-05 17:42:51]: Step 8 of 20 Updating recovery media
[RSNG 2012-06-05 17:42:51]: Step 9 of 20 Extracting incoming image
[RSNG 2012-06-05 17:42:51]: Step 10 of 20 Unpacking OS packages
[RSNG 2012-06-05 17:42:51]: Step 11 of 20 Mounting jbase package
[RSNG 2012-06-05 17:42:51]: Step 12 of 20 Creating base OS symbolic links
[RSNG 2012-06-05 17:42:51]: Step 13 of 20 Creating fstab
[RSNG 2012-06-05 17:42:51]: Step 14 of 20 Creating new system files
[RSNG 2012-06-05 17:42:51]: Step 15 of 20 Adding jbundle package
[RSNG 2012-06-05 17:42:51]: Step 16 of 20 Backing up system data
[RSNG 2012-06-05 17:42:51]: Step 17 of 20 Setting up shared partition data
[RSNG 2012-06-05 17:42:51]: Step 18 of 20 Checking package sanity in
installation
[RSNG 2012-06-05 17:42:51]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[RSNG 2012-06-05 17:42:51]: Step 20 of 20 Setting da0s2 as new active partition
[RSNG 2012-06-05 17:43:36]: Rebooting Backup RE
[RSNG 2012-06-05 17:43:36]: ----- Rebooting fpc1 -----
[RSNG 2012-06-05 17:50:12]: Initiating Chassis In-Service-Upgrade
[RSNG 2012-06-05 17:50:33]: Upgrading group: 0 fpc: 0
[RSNG 2012-06-05 17:52:38]: Upgrade complete for group:0
[RSNG 2012-06-05 17:52:38]: Upgrading group: 1 fpc: 1
[RSNG 2012-06-05 17:54:42]: Upgrade complete for group:1
[RSNG 2012-06-05 17:54:42]: Finished processing all upgrade groups, last group
:1
[RSNG 2012-06-05 17:54:48]: Preparing for Switchover
[RSNG 2012-06-05 17:55:38]: Switchover Completed
```



```
[Status 2012-06-05 17:55:41]: Upgrade completed with 0 errors
Success
```

request system software nonstop-upgrade node-group (Server Node Group)

```
user@qfabric> request system software nonstop-upgrade node-group P1507-C
jinstall-qfabric-12.2X50-D10.3.rpm
Upgrading target(s): P1507-C

[P1507-C 2012-06-26 14:02:44]: Retrieving package
[P1507-C 2012-06-26 14:03:21]: ----- P1507-C: -----
[P1507-C 2012-06-26 14:03:59]: Step 1 of 20 Creating temporary file system
[P1507-C 2012-06-26 14:03:59]: Step 2 of 20 Determining installation source
[P1507-C 2012-06-26 14:03:59]: Step 3 of 20 Processing format options
[P1507-C 2012-06-26 14:03:59]: Step 4 of 20 Determining installation slice
[P1507-C 2012-06-26 14:04:00]: Step 5 of 20 Creating and labeling new slices
[P1507-C 2012-06-26 14:04:00]: Step 6 of 20 Create and mount new file system
[P1507-C 2012-06-26 14:04:08]: Step 7 of 20 Getting OS bundles
[P1507-C 2012-06-26 14:04:09]: Step 8 of 20 Updating recovery media
[P1507-C 2012-06-26 14:04:29]: Step 9 of 20 Extracting incoming image
[P1507-C 2012-06-26 14:05:42]: Step 10 of 20 Unpacking OS packages
[P1507-C 2012-06-26 14:05:49]: Step 11 of 20 Mounting jbase package
[P1507-C 2012-06-26 14:06:14]: Step 12 of 20 Creating base OS symbolic links
[P1507-C 2012-06-26 14:07:15]: Step 13 of 20 Creating fstab
[P1507-C 2012-06-26 14:07:15]: Step 14 of 20 Creating new system files
[P1507-C 2012-06-26 14:07:16]: Step 15 of 20 Adding jbundle package
[P1507-C 2012-06-26 14:09:52]: Step 16 of 20 Backing up system data
[P1507-C 2012-06-26 14:10:07]: Step 17 of 20 Setting up shared partition data
[P1507-C 2012-06-26 14:10:07]: Step 18 of 20 Checking package sanity in
installation
[P1507-C 2012-06-26 14:10:08]: Step 19 of 20 Unmounting and cleaning up temporary
file systems
[P1507-C 2012-06-26 14:10:11]: Step 20 of 20 Setting da0s2 as new active partition
[Status 2012-06-26 14:10:25]: Trying to Connect to Node: P1507-C
[Status 2012-06-26 14:10:32]: Rebooting P1507-C
[Status 2012-06-26 14:10:32]: Upgrade completed with 0 errors
Success
```

show chassis nonstop-upgrade node-group

Syntax	show chassis nonstop-upgrade node-group <i>node-group-name</i>
Release Information	Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Display the status of the Node group after the most recent nonstop software upgrade (NSSU).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Performing a Nonstop Software Upgrade on the QFabric System on page 36 • request system software nonstop-upgrade on page 164
List of Sample Output	show chassis nonstop-upgrade node-group on page 170
Output Fields	Table 15 on page 170 lists the output fields for the show chassis nonstop-upgrade node-group command. Output fields are listed in the approximate order in which they appear.

Table 15: show chassis nonstop-upgrade node-group Output Fields

Field Name	Field Description
Item	Node device slot number.
Status	State of Node device: <ul style="list-style-type: none"> • Error—Node device is in an error state. • Offline—Node device is powered down. • Online—Node device is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade node-group

```

user@qfabric> show chassis nonstop-upgrade node-group NW-NG-0
Item           Status           Reason
P1550-C        Online

```

CHAPTER 11

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 177</p> <p>show vrrp brief on page 177</p> <p>show vrrp detail (IPv6) on page 177</p> <p>show vrrp detail (Route Track) on page 178</p> <p>show vrrp extensive on page 178</p> <p>show vrrp interface on page 179</p> <p>show vrrp summary on page 180</p> <p>show vrrp track detail on page 180</p> <p>show vrrp track summary on page 181</p>
Output Fields	Table 16 on page 172 lists the output fields for the show vrrp command. Output fields are listed in the approximate order in which they appear.

Table 16: show vrrp Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the logical interface.	none, brief, extensive, summary
Interface index	Physical interface index number, which reflects its initialization sequence.	extensive
Groups	Total number of VRRP groups configured on the interface.	extensive

Table 16: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	extensive
Interface VRRP PDU statistics	Nonerrored statistics for the logical interface: <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 16: 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 16: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Master router	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is N/A .	detail, extensive
Virtual router uptime	Time that the virtual router has been up.	detail, extensive
Master router uptime	Time that the master router has been up.	detail, extensive
Virtual MAC	MAC address associated with the virtual IP address.	detail, extensive
Tracking	Whether tracking is enabled or disabled .	detail, extensive
Current priority	Current operational priority for being the VRRP master.	detail, extensive
Configured priority	Configured base priority for being the VRRP master.	detail, extensive
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. Disabled indicates no minimum interval.	detail, extensive
Remaining-time	(track option only) Displays the time remaining in the priority hold-time interval.	detail
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	detail extensive
Interface/Tracked interface	Name of the tracked interface.	detail extensive
Int state/Interface state	Current operational state of the tracked interface: up or down .	detail, extensive
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	detail, extensive
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	detail, extensive
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of down means that the corresponding priority cost is incurred when the interface is down.	detail, extensive
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	detail, extensive
Route count	The number of routes being tracked.	detail, extensive

Table 16: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Route	The IP address of the route being tracked.	detail, extensive
VRF name	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	detail, extensive
Route state	The state of the route being tracked: up , down , or unknown .	detail, extensive
Priority cost	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	detail, extensive
Active	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	detail, extensive
Group VRRP PDU statistics	Number of VRRP advertisements sent and received by the group.	extensive
Group VRRP PDU error statistics	<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 16: 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  gec0::12:1:1:1
                                     vip  ge80::12:1:1:99
                                     vip  gec0::12:1:1:99
ge-0/0/2.131   up         1      master    A 0.364 1c1  gec0::13:1:1:1
                                     vip  ge80::13:1:1:99
                                     vip  gec0::13:1:1:99

```

show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see [show vrrp on page 177](#).

show vrrp detail (IPv6)

```

user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
 Interface state: up, Group: 1, State: master
 Priority: 200, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
 gec0::13:1:1:99
 Advertisement timer: 0.327s, Master router: ge80::13:1:1:1
 Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
 Virtual MAC: 00:00:5e:00:02:01
 Tracking: disabled

show vrrp detail (Route Track)

user@host> show vrrp detail

Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24
 Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled
 Interface state: up, Group: 100, State: master
 Priority: 150, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
 Advertisement timer: 1.218s, Master router: 30.30.30.30
 Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
 Virtual MAC: 00:00:5e:00:01:64
 Tracking: enabled
 Current priority: 150, Configured priority: 150
 Priority hold-time: disabled
 Interface tracking: disabled
 Route tracking: enabled, Route count: 1

Route	VRF name	Route state	Priority cost
192.168.40.0/22	default	up	30

show vrrp extensive

user@host> show vrrp extensive

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1

Interface VRRP PDU statistics

Advertisement sent	:	188
Advertisement received	:	0
Packets received	:	0
No group match received	:	0

Interface VRRP PDU error statistics

Invalid IPAH next type received	:	0
Invalid VRRP TTL value received	:	0
Invalid VRRP version received	:	0
Invalid VRRP PDU type received	:	0
Invalid VRRP authentication type received	:	0
Invalid VRRP IP count received	:	0
Invalid VRRP checksum received	:	0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
 Interface state: up, Group: 1, State: master
 Priority: 200, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
 gec0::12:1:1:99
 Advertisement timer: 1.034s, Master router: ge80::12:1:1:1
 Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
 Virtual MAC: 00:00:5e:00:02:01
 Tracking: disabled
 Group VRRP PDU statistics

```

    Advertisement sent          :          188
    Advertisement received      :           0
Group VRRP PDU error statistics
    Bad authentication type received:          0
    Bad password received       :           0
    Bad MD5 digest received     :           0
    Bad advertisement timer received:          0
    Bad VIP count received      :           0
    Bad VIPADDR received       :           0
Group state transition statistics
    Idle to master transitions   :           0
    Idle to backup transitions  :           1
    Backup to master transitions :           1
    Master to backup transitions :           0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
Interface VRRP PDU statistics
    Advertisement sent          :          186
    Advertisement received      :           0
    Packets received            :           0
    No group match received     :           0
Interface VRRP PDU error statistics
    Invalid IPAH next type received :          0
    Invalid VRRP TTL value received :          0
    Invalid VRRP version received  :          0
    Invalid VRRP PDU type received :          0
    Invalid VRRP authentication type received:          0
    Invalid VRRP IP count received :          0
    Invalid VRRP checksum received :          0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
Advertisement timer: 0.396s, Master router: ge80::13:1:1:1
Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
    Advertisement sent          :          186
    Advertisement received      :           0
Group VRRP PDU error statistics
    Bad authentication type received:          0
    Bad password received       :           0
    Bad MD5 digest received     :           0
    Bad advertisement timer received:          0
    Bad VIP count received      :           0
    Bad VIPADDR received       :           0
Group state transition statistics
    Idle to master transitions   :           0
    Idle to backup transitions  :           1
    Backup to master transitions :           1
    Master to backup transitions :           0

```

show vrrp interface

user@host> show vrrp interface

```

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
Interface VRRP PDU statistics
  Advertisement sent           :          205
  Advertisement received       :           0
  Packets received             :           0
  No group match received      :           0
Interface VRRP PDU error statistics
  Invalid IPAH next type received :           0
  Invalid VRRP TTL value received :           0
  Invalid VRRP version received  :           0
  Invalid VRRP PDU type received :           0
  Invalid VRRP authentication type received:           0
  Invalid VRRP IP count received :           0
  Invalid VRRP checksum received :           0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 0.789s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent           :          205
  Advertisement received       :           0
Group VRRP PDU error statistics
  Bad authentication type received:           0
  Bad password received         :           0
  Bad MD5 digest received       :           0
  Bad advertisement timer received:           0
  Bad VIP count received        :           0
  Bad VIPADDR received          :           0
Group state transition statistics
  Idle to master transitions     :           0
  Idle to backup transitions     :           1
  Backup to master transitions   :           1
  Master to backup transitions   :           0

```

show vrrp summary

```

user@host> show vrrp summary
Interface   State   Group  VR state  Type  Address
ge-4/1/0.0  up      1      backup    lcl   10.57.0.2
vip         vip     vip     vip     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

