



Junos[®] OS for EX Series Ethernet Switches

High Availability for EX9200 Switches

Release

12.3



Published: 2013-04-01

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

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS for EX Series Ethernet Switches High Availability for EX9200 Switches

Release 12.3

Copyright © 2013, 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	xiii
	Documentation and Release Notes	xiii
	Supported Platforms	xiii
	Using the Examples in This Manual	xiii
	Merging a Full Example	xiv
	Merging a Snippet	xiv
	Documentation Conventions	xv
	Documentation Feedback	xvii
	Requesting Technical Support	xvii
	Self-Help Online Tools and Resources	xvii
	Opening a Case with JTAC	xviii
Part 1	Overview	
Chapter 1	Graceful Routing Engine Switchover (GRES)	3
	Understanding Graceful Routing Engine Switchover in the Junos OS	3
	Graceful Routing Engine Switchover Concepts	3
	Effects of a Routing Engine Switchover	6
	Graceful Routing Engine Switchover System Requirements	7
	Graceful Routing Engine Switchover Platform Support	7
	Graceful Routing Engine Switchover Feature Support	7
	Graceful Routing Engine Switchover DPC Support	9
	Graceful Routing Engine Switchover and Subscriber Access	9
	Graceful Routing Engine Switchover PIC Support	9
	Requirements for Routers with a Backup Router Configuration	10
Chapter 2	Nonstop Bridging (NSB)	11
	Nonstop Bridging Concepts	11
	Nonstop Bridging System Requirements	13
	Platform Support	13
	Protocol Support	14
Chapter 3	Nonstop Active Routing (NSR)	15
	Nonstop Active Routing Concepts	15
	Nonstop Active Routing System Requirements	18
	Nonstop Active Routing Platform Support	18
	Nonstop Active Routing Protocol and Feature Support	19
	Nonstop Active Routing BFD Support	21
	Nonstop Active Routing BGP Support	22
	Nonstop Active Routing Layer 2 Circuit and VPLS Support	23
	Nonstop Active Routing PIM Support	23
	Nonstop Active Routing MSDP Support	26

	Nonstop Active Routing Support for RSVP-TE LSPs	27
Chapter 4	Graceful Restart	29
	Graceful Restart Concepts	29
	Graceful Restart System Requirements	30
	Aggregate and Static Routes	31
	Graceful Restart and Routing Protocols	31
	BGP	31
	ES-IS	32
	IS-IS	32
	OSPF and OSPFv3	32
	PIM Sparse Mode	33
	RIP and RIPng	33
	Graceful Restart and MPLS-Related Protocols	34
	LDP	34
	RSVP	34
	CCC and TCC	35
	Graceful Restart and Layer 2 and Layer 3 VPNs	35
	Graceful Restart on Logical Systems	36
Chapter 5	Unified ISSU	37
	Upgrading Routers Using ISSU	37
	Unified ISSU Concepts	37
	Unified ISSU System Requirements	42
	Unified ISSU Junos OS Release Support	42
	Unified ISSU Platform Support	43
	Unified ISSU Protocol Support	43
	Unified ISSU Support for the Layer 2 Control Protocol Process	44
	Unified ISSU Feature Support	45
	Unified ISSU PIC Support	45
	PIC Considerations	46
	SONET/SDH PICs	46
	Fast Ethernet and Gigabit Ethernet PICs	48
	Channelized PICs	49
	Tunnel Services PICs	50
	ATM PICs	50
	Serial PICs	51
	DS3, E1, E3, and T1 PICs	51
	Enhanced IQ PICs	52
	Enhanced IQ2 Ethernet Services Engine (ESE) PIC	52
	Unified ISSU Support on MX Series 3D Universal Edge Routers	53
	Unified ISSU DPC and FPC Support on MX Series 3D Universal Edge Routers	53
	Unified ISSU MIC and MPC Support on MX Series 3D Universal Edge Routers	53
	Unified ISSU Limitation on MX Series 3D Universal Edge Routers	54

Chapter 6	VRRP	55
	Understanding VRRP	55
	Junos OS Support for VRRPv3	56
	Understanding VRRPv3 Behavioral Differences	57
	Understanding VRRPv2 to VRRPv3 Transition	58
	Improving the Convergence Time for VRRP	59
Part 2	Configuration	
Chapter 7	Configuration: GRES	65
	Configuring Graceful Routing Engine Switchover	65
	Enabling Graceful Routing Engine Switchover	65
	Synchronizing the Routing Engine Configuration	65
	Verifying Graceful Routing Engine Switchover Operation	66
	Resetting Local Statistics	66
Chapter 8	Configuration Statements: GRES	69
	[edit chassis] Hierarchy Level	69
	graceful-switchover	77
Chapter 9	Configuration: Graceful Restart	79
	Enabling Graceful Restart	79
	Configuring Routing Protocols Graceful Restart	80
	Enabling Graceful Restart	80
	Configuring Graceful Restart Options for BGP	81
	Configuring Graceful Restart Options for ES-IS	82
	Configuring Graceful Restart Options for IS-IS	82
	Configuring Graceful Restart Options for OSPF and OSPFv3	83
	Configuring Graceful Restart Options for RIP and RIPng	84
	Configuring Graceful Restart Options for PIM Sparse Mode	84
	Tracking Graceful Restart Events	86
	Configuring Graceful Restart for MPLS-Related Protocols	86
	Configuring Graceful Restart Globally	86
	Configuring Graceful Restart Options for RSVP, CCC, and TCC	87
	Configuring Graceful Restart Options for LDP	87
	Configuring VPN Graceful Restart	88
	Configuring Graceful Restart Globally	88
	Configuring Graceful Restart for the Routing Instance	89
	Configuring Logical System Graceful Restart	89
	Enabling Graceful Restart Globally	90
	Configuring Graceful Restart for a Routing Instance	90
	Example: Configuring Graceful Restart	91
	Example: Managing Helper Modes for OSPF Graceful Restart	116
Chapter 10	Configuration Statements: Graceful Restart	119
	disable	119
	graceful-restart (Enabling Globally)	120
	helper-disable (Multiple Protocols)	121
	maximum-helper-recovery-time	121
	maximum-helper-restart-time (RSVP)	122

	maximum-neighbor-reconnect-time	122
	maximum-neighbor-recovery-time	123
	no-strict-lsa-checking	124
	notify-duration	125
	reconnect-time	126
	recovery-time	127
	restart-duration	128
	restart-time (BGP Graceful Restart)	129
	stale-routes-time	130
	traceoptions (Protocols)	131
Chapter 11	Configuration: NSB	133
	Configuring Nonstop Bridging	133
	Enabling Nonstop Bridging	133
	Synchronizing the Routing Engine Configuration	133
	Verifying Nonstop Bridging Operation	134
	Resetting Local Statistics	134
Chapter 12	Configuration Statements: NSB	135
	[edit protocols layer2-control] Hierarchy Level	135
	nonstop-bridging	136
Chapter 13	Configuration: NSR	137
	Configuring Nonstop Active Routing	137
	Enabling Nonstop Active Routing	137
	Synchronizing the Routing Engine Configuration	138
	Verifying Nonstop Active Routing Operation	138
	Tracing Nonstop Active Routing Synchronization Events	139
	traceoptions (Routing Options)	141
	Example: Configuring Nonstop Active Routing	143
Chapter 14	Configuration Statements: NSR	147
	[edit protocols layer2-control] Hierarchy Level	147
	commit synchronize	148
	nonstop-routing	149
	traceoptions (Routing Options)	150
Chapter 15	Configuration: Unified ISSU	153
	Best Practices	153
	Before You Begin	154
	Verify That the Master and Backup Routing Engines Are Running the Same Software Version	155
	Back Up the Router Software	155
	Verify That Graceful Routing Engine Switchover and Nonstop Active Routing Are Configured	156
	Performing a Unified ISSU	157
	Upgrading and Rebooting Both Routing Engines Automatically	157
	Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually	161

	Upgrading and Rebooting Only One Routing Engine	166
	Verifying a Unified ISSU	169
	Managing and Tracing BFD Sessions During Unified ISSU Procedures	170
Chapter 16	Configuration Statements: Unified ISSU	173
	bfd	174
	no-issu-timer-negotiation	176
	traceoptions (Protocols BFD)	177
Chapter 17	Configuration: VRRP	179
	Configuring the Startup Period for VRRP Operations	179
	Configuring Basic VRRP Support	180
	Configuring VRRP Authentication (IPv4 Only)	182
	Configuring the Advertisement Interval for the VRRP Master Router	184
	Modifying the Advertisement Interval in Seconds	184
	Modifying the Advertisement Interval in Milliseconds	185
	Configuring a Backup Router to Preempt the Master Router	186
	Modifying the Preemption Hold-Time Value	187
	Configuring Asymmetric Hold Time for VRRP Routers	187
	Configuring an Interface to Accept Packets Destined for the Virtual IP Address	188
	Configuring a Logical Interface to Be Tracked	189
	Configuring a Route to Be Tracked	191
	Configuring Inheritance for a VRRP Group	192
	Configuring the Silent Period	193
	Configuring Passive ARP Learning for Backup VRRP Routers	194
	Enabling the Distributed Periodic Packet Management Process for VRRP	195
	Configuring VRRP to Improve Convergence Time	196
	Example: Configuring VRRP	197
	Example: Configuring VRRP for IPv6	199
	Example: Configuring VRRP Route Tracking	200
	Tracing VRRP Operations	201
Chapter 18	Configuration Statements: VRRP	203
	[edit protocols vrrp] Hierarchy Level	203
	accept-data	204
	advertise-interval	205
	asymmetric-hold-time	206
	authentication-key	207
	authentication-type	208
	bandwidth-threshold	209
	delegate-processing (VRRP)	210
	failover-delay	210
	fast-interval	211
	global-advertisements-threshold	212
	hold-time (VRRP)	213
	inet6-advertise-interval	214
	interface (VRRP Group)	215
	preempt (VRRP)	216
	priority (Protocols VRRP)	217

	priority-cost (VRRP)	218
	priority-hold-time	219
	route (Interfaces)	220
	skew-timer-disable	221
	startup-silent-period	221
	traceoptions (Protocols VRRP)	222
	track (VRRP)	224
	version-3	225
	virtual-address	226
	virtual-inet6-address	226
	virtual-link-local-address	227
	vrrp-group	228
	vrrp-inet6-group	229
Part 3	Administration	
Chapter 19	Routine Monitoring	233
	Resetting Local Statistics	233
	Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart	234
	Verifying Graceful Restart Operation	234
	Graceful Restart Operational Mode Commands	235
	Verifying BGP Graceful Restart	235
	Verifying IS-IS and OSPF Graceful Restart	236
	Verifying CCC and TCC Graceful Restart	236
Chapter 20	Operational Commands	239
	show system switchover	240
Part 4	Troubleshooting	
Chapter 21	Troubleshooting Unified ISSU	247
	Troubleshooting Unified ISSU Problems	247

List of Figures

Part 1	Overview	
Chapter 1	Graceful Routing Engine Switchover (GRES)	3
	Figure 1: Preparing for a Graceful Routing Engine Switchover	4
	Figure 2: Graceful Routing Engine Switchover Process	5
Chapter 2	Nonstop Bridging (NSB)	11
	Figure 3: Nonstop Bridging Switchover Preparation Process	12
	Figure 4: Nonstop Bridging During a Switchover	13
Chapter 3	Nonstop Active Routing (NSR)	15
	Figure 5: Nonstop Active Routing Switchover Preparation Process	16
	Figure 6: Nonstop Active Routing During a Switchover	17
Chapter 6	VRRP	55
	Figure 7: Basic VRRP	56
Part 2	Configuration	
Chapter 9	Configuration: Graceful Restart	79
	Figure 8: Layer 3 VPN Graceful Restart Topology	91

List of Tables

	About the Documentation	xiii
	Table 1: Notice Icons	xv
	Table 2: Text and Syntax Conventions	xv
Part 1	Overview	
Chapter 1	Graceful Routing Engine Switchover (GRES)	3
	Table 3: Effects of a Routing Engine Switchover	6
	Table 4: Graceful Routing Engine Switchover Feature Support	7
Chapter 3	Nonstop Active Routing (NSR)	15
	Table 5: Nonstop Active Routing Platform Support	18
	Table 6: Nonstop Active Routing Protocol and Feature Support	19
Chapter 5	Unified ISSU	37
	Table 7: Unified ISSU Platform Support	43
	Table 8: Unified ISSU Protocol Support	43
	Table 9: Unified ISSU PIC Support: SONET/SDH	47
	Table 10: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet	48
	Table 11: Unified ISSU PIC Support: Channelized	50
	Table 12: Unified ISSU PIC Support: Tunnel Services	50
	Table 13: Unified ISSU PIC Support: ATM	51
	Table 14: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC	52
	Table 15: Unified ISSU Support: MX Series 3D Universal Edge Routers	53
	Table 16: Unified ISSU Support: MX Series 3D Universal Edge Routers	54
Chapter 6	VRRP	55
	Table 17: Example: VRRPv2 to VRRPv3 Transition Steps and Events	58
Part 2	Configuration	
Chapter 17	Configuration: VRRP	179
	Table 18: Interface State and Priority Cost Usage	190
Part 3	Administration	
Chapter 20	Operational Commands	239
	Table 19: show system switchover Output Fields	242

About the Documentation

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

Documentation and Release Notes

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

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

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

Supported Platforms

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

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active

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

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book 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 System Basics Configuration 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)	Enclose optional keywords or variables.	stub <default-metric metric>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast <i>(string1 string2 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)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identify 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.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web 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.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>

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

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

Opening a Case with JTAC

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

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Overview

- [Graceful Routing Engine Switchover \(GRES\) on page 3](#)
- [Nonstop Bridging \(NSB\) on page 11](#)
- [Nonstop Active Routing \(NSR\) on page 15](#)
- [Graceful Restart on page 29](#)
- [Unified ISSU on page 37](#)
- [VRRP on page 55](#)

CHAPTER 1

Graceful Routing Engine Switchover (GRES)

- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 3](#)
- [Graceful Routing Engine Switchover System Requirements on page 7](#)
- [Requirements for Routers with a Backup Router Configuration on page 10](#)

Understanding Graceful Routing Engine Switchover in the Junos OS

This topic contains the following sections:

- [Graceful Routing Engine Switchover Concepts on page 3](#)
- [Effects of a Routing Engine Switchover on page 6](#)

Graceful Routing Engine Switchover Concepts

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. Graceful Routing Engine switchover preserves interface and kernel information. Traffic is not interrupted. However, graceful Routing Engine switchover does not preserve the control plane. 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, graceful Routing Engine switchover must be combined with either graceful restart protocol extensions or nonstop active routing. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.



NOTE: To quickly restore or to preserve routing protocol state information during a switchover, graceful Routing Engine switchover must be combined with either graceful restart or nonstop active routing (NSR), respectively. For more information about graceful restart, see [“Graceful Restart Concepts” on page 29](#). For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 15](#).

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 and reconnects to the new master Routing Engine. The Packet Forwarding Engine does not reboot, and traffic is not interrupted. 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: Successive Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

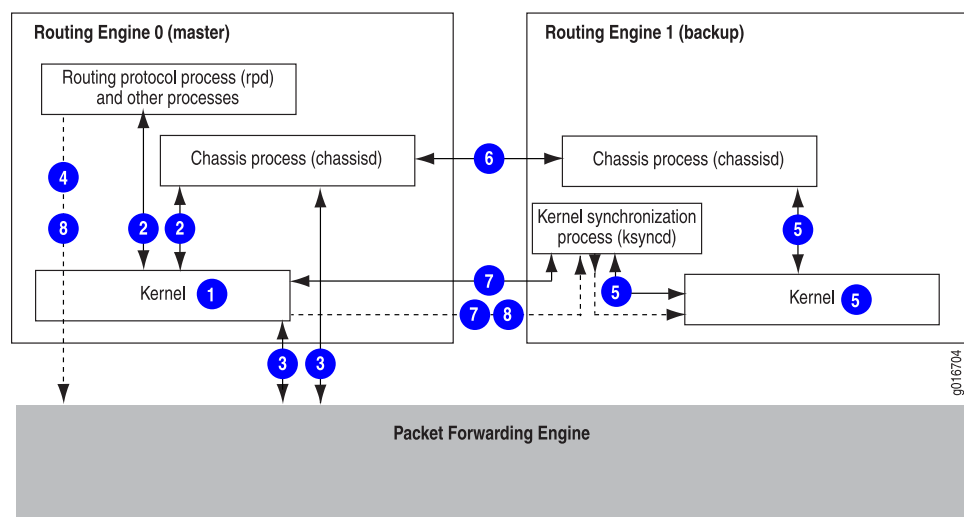
If the router displays a warning message similar to “Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset,” do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines 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: We do not recommend performing a commit operation on the backup Routing Engine when graceful Routing Engine switchover is enabled on the router.

Figure 1 on page 4 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

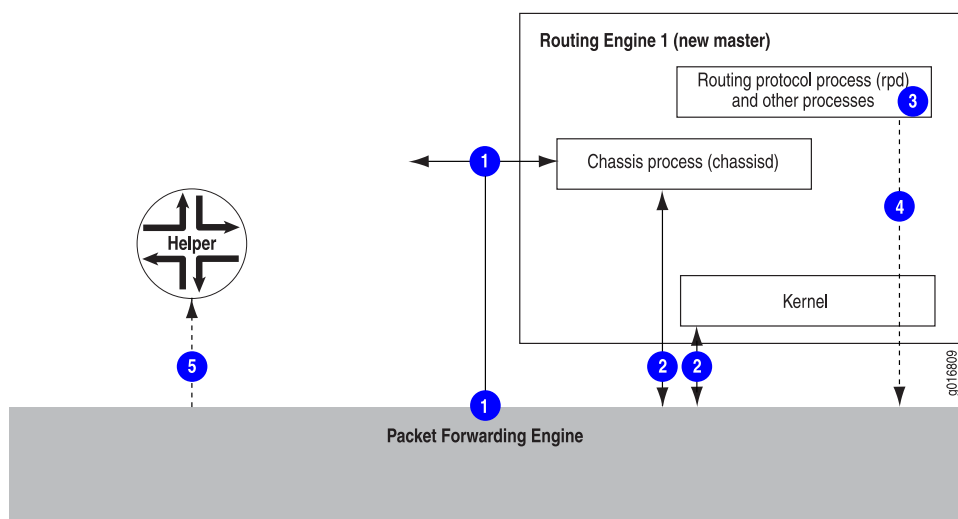


The switchover preparation process for graceful Routing Engine switchover follows these steps:

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 graceful Routing Engine switchover 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 5 shows the effects of a switchover on the routing platform.

Figure 2: Graceful Routing Engine Switchover Process



When a switchover occurs, the switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master.
3. Routing platform processes that are not part of graceful Routing Engine switchover (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, the Switch Interface Boards (SIBs) are taken offline and restarted one by one during a graceful Routing Engine switchover. This is done to provide the 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.

Effects of a Routing Engine Switchover

Table 3 on page 6 describes the effects of a Routing Engine switchover when no high availability features are enabled and when graceful Routing Engine switchover, graceful restart, and nonstop active routing features are enabled.

Table 3: Effects of a Routing Engine Switchover

Feature	Benefits	Considerations
Dual Routing Engines only (no features enabled)	When the switchover to the new master Routing Engine is complete, routing convergence takes place and traffic is resumed.	All physical interfaces are taken offline, Packet Forwarding Engines restart, the standby Routing Engine restarts the routing protocol process (rpd), and all hardware and interfaces are discovered by the new master Routing Engine. The switchover takes several minutes and all of the router's adjacencies are aware of the physical (interface alarms) and routing (topology) change.
Graceful Routing Engine switchover enabled	During the switchover, interface and kernel information is preserved. The switchover is faster because the Packet Forwarding Engines are not restarted.	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.
Graceful Routing Engine switchover and nonstop active routing enabled	Traffic is not interrupted during the switchover. Interface, kernel, and routing protocol information is preserved.	Unsupported protocols must be refreshed using the normal recovery mechanisms inherent in each protocol.
Graceful Routing Engine switchover and graceful restart enabled	Traffic is not interrupted during the switchover. Interface and kernel information is preserved. Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers.	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.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Graceful Routing Engine Switchover System Requirements on page 7](#)
- [Configuring Graceful Routing Engine Switchover on page 65](#)
- [Requirements for Routers with a Backup Router Configuration on page 10](#)

Graceful Routing Engine Switchover System Requirements

Graceful Routing Engine switchover is supported on all routing 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 7](#)
- [Graceful Routing Engine Switchover Feature Support on page 7](#)
- [Graceful Routing Engine Switchover DPC Support on page 9](#)
- [Graceful Routing Engine Switchover and Subscriber Access on page 9](#)
- [Graceful Routing Engine Switchover PIC Support on page 9](#)

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
- T1600 router—Junos OS Release 8.5 or later
- T4000 router—Junos OS Release 12.1R2 or later
- TX Matrix Plus router—Junos OS Release 9.6 or later

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 7](#).

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

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

Application	Junos OS Release
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.

- VRRP changes mastership when a Routing Engine switchover occurs, even when graceful Routing Engine switchover is configured.

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. 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 in the Junos OS on page 3](#)
 - [Configuring Graceful Routing Engine Switchover on page 65](#)
 - [Requirements for Routers with a Backup Router Configuration on page 10](#)

Requirements for Routers with a Backup Router Configuration

If your Routing Engine configuration includes a **backup-router** statement or an **inet6-backup-router** statement, you can also use the **destination** statement to specify a subnet address or multiple subnet addresses for the backup router. Include destination subnets for the backup Routing Engine at the **[edit system (backup-router | inet6-backup-router) address]** hierarchy level. This requirement also applies to any T640 router connected to a TX Matrix router that includes a **backup-router** or **inet6-backup-router** statement.



NOTE: If you have a backup router configuration in which multiple static routes point to a gateway from the management Ethernet interface, you must configure prefixes that are more specific than the static routes or include the retain flag at the **[edit routing-options static route]** hierarchy level.

For example, if you configure the static route 172.16.0.0/12 from the management Ethernet interface for management purposes, you must specify the backup router configuration as follows:

```
backup-router 172.29.201.62 destination [172.16.0.0/13 172.16.128.0/13]
```

- Related Documentation**
- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 3](#)
 - [Graceful Routing Engine Switchover System Requirements on page 7](#)

CHAPTER 2

Nonstop Bridging (NSB)

- [Nonstop Bridging Concepts on page 11](#)
- [Nonstop Bridging System Requirements on page 13](#)

Nonstop Bridging Concepts


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



NOTE: To use nonstop bridging, you must first enable graceful Routing Engine switchover on your routing platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover in the Junos OS” on page 3](#).

Figure 3 on page 12 shows the system architecture of nonstop bridging and the process a routing platform follows to prepare for a switchover.

Figure 3: Nonstop Bridging Switchover Preparation Process



The switchover preparation process for nonstop bridging follows these steps:

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

Figure 4 on page 13 shows the effects of a switchover on the routing platform.



Nonstop bridging is supported on EX Series Ethernet Switch with redundant Routing Engines.

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



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

Protocol Support

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

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

Related Documentation

- [Nonstop Bridging Concepts on page 11](#)
- [Configuring Nonstop Bridging on page 133](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)

CHAPTER 3

Nonstop Active Routing (NSR)

- [Nonstop Active Routing Concepts on page 15](#)
- [Nonstop Active Routing System Requirements on page 18](#)

Nonstop Active Routing Concepts


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



NOTE: To use nonstop active routing, you must first enable graceful Routing Engine switchover on your routing platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover in the Junos OS” on page 3](#).

Figure 5 on page 16 shows the system architecture of nonstop active routing and the process a routing platform follows to prepare for a switchover.

Figure 5: Nonstop Active Routing Switchover Preparation Process




The switchover preparation process for nonstop active routing follows these steps:

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

Figure 6 on page 17 shows the effects of a switchover on the routing platform.

Figure 6: Nonstop Active Routing During a Switchover



The switchover process follows these steps:

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

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Active Routing System Requirements on page 18](#)
- [Configuring Nonstop Active Routing on page 137](#)

Nonstop Active Routing System Requirements

This section contains the following topics:

- [Nonstop Active Routing Platform Support on page 18](#)
- [Nonstop Active Routing Protocol and Feature Support on page 19](#)
- [Nonstop Active Routing BFD Support on page 21](#)
- [Nonstop Active Routing BGP Support on page 22](#)
- [Nonstop Active Routing Layer 2 Circuit and VPLS Support on page 23](#)
- [Nonstop Active Routing PIM Support on page 23](#)
- [Nonstop Active Routing MSDP Support on page 26](#)
- [Nonstop Active Routing Support for RSVP-TE LSPs on page 27](#)

Nonstop Active Routing Platform Support

[Table 5 on page 18](#) lists the platforms that support nonstop active routing (NSR).

Table 5: Nonstop Active Routing Platform Support

Platform	Junos OS Release
M10i router	8.4 or later
M20 router	8.4 or later
M40e router	8.4 or later
M120 router	9.0 or later
M320 router	8.4 or later
MX Series routers	9.0 or later
PTX Series Packet Transport switches	12.1R4 or later
<p>NOTE:</p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> • Labeled BGP • Layer 2 VPNs excluding Layer 2 interworking (Layer 2 stitching) • Layer 3 VPNs • LDP • RSVP 	
T320 router, T640 router, and TX Matrix router	8.4 or later
T1600 router	8.5 or later

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

Platform	Junos OS Release
TX Plus Matrix router	10.0 or later



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

Nonstop Active Routing Protocol and Feature Support

Table 6 on page 19 lists the protocols that are supported by nonstop active routing.

Table 6: Nonstop Active Routing Protocol and Feature Support

Protocol	Junos OS Release
Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP)	9.4 or later
Bidirectional Forwarding Detection (BFD)	8.5 or later
For more information, see “Nonstop Active Routing BFD Support” on page 21.	
BGP	8.4 or later
For more information, see “Nonstop Active Routing BGP Support” on page 22.	
Labeled BGP (PTX Series Packet Transport Switches only)	12.1R4 or later
IS-IS	8.4 or later
LDP	8.4 or later
LDP-based virtual private LAN service (VPLS)	9.3 or later
LDP OAM (operation, administration, and management) features	9.6 or later

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

Protocol	Junos OS Release
LDP (PTX Series Packet Transport Switches only)	12.1R4 or later
Nonstop active routing support for LDP includes:	
<ul style="list-style-type: none"> • LDP unicast transit LSPs • LDP egress LSPs for labeled internal BGP (IBGP) and external BGP (EBGP) • LDP over RSVP transit LSPs • LDP transit LSPs with indexed next hops • LDP transit LSPs with unequal cost load balancing 	
NOTE: Nonstop active routing is not supported for LDP Point-to-Multipoint LSPs and LDP ingress LSPs.	
Layer 2 circuits	(on LDP-based VPLS) 9.2 or later (on RSVP-TE LSP) 11.1 or later
Layer 2 VPNs	9.1 or later
Layer 2 VPNs (PTX Series Packet Transport Switches only)	12.1R4 or later
NOTE: Nonstop active routing is not supported for Layer 2 interworking (Layer 2 stitching).	
Layer 3 VPNs (see the first Note after this table for restrictions)	9.2 or later
Layer 3 VPNs (PTX Series Packet Transport Switches only)	12.1R4 or later
Multicast Source Discovery Protocol (MSDP)	12.1 or later
For more information, see “Nonstop Active Routing MSDP Support” on page 26 .	
OSPF/OSPFv3	8.4 or later
Protocol Independent Multicast (PIM)	(for IPv4) 9.3 or later
For more information, see “Nonstop Active Routing PIM Support” on page 23 .	(for IPv6) 10.4 or later
RIP and RIP next generation (RIPng)	9.0 or later

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

Protocol	Junos OS Release
RSVP (PTX Series Packet Transport Switches only)	12.1R4 or later
Nonstop active routing support for RSVP includes:	
<ul style="list-style-type: none"> Point-to-Multipoint LSPs <ul style="list-style-type: none"> RSVP Point-to-Multipoint ingress, transit, and egress LSPs using existing non-chained next hop. RSVP Point-to-Multipoint transit LSPs using composite next hops for Point-to-Multipoint label routes. Point-to-Point LSPs <ul style="list-style-type: none"> RSVP Point-to-Point ingress, transit, and egress LSPs using non-chained next hops. RSVP Point-to-Point transit LSPs using chained composite next hops. 	
RSVP-TE LSP	9.5 or later
For more information, see “Nonstop Active Routing Support for RSVP-TE LSPs” on page 27 .	
VPLS	(LDP-based) 9.1 or later (RSVP-TE-based) 11.2 or later



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



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



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

Nonstop Active Routing BFD Support

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



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

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



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

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

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions, and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

Nonstop Active Routing BGP Support

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

- You must include the **path-selection external-router-ID** statement at the **[edit protocols bgp]** hierarchy level to ensure consistent path selection between the master and backup Routing Engines during and after the nonstop active routing switchover.
- If the BGP peer in the master Routing Engine has negotiated address-family capabilities that are not supported for nonstop active routing, then the corresponding BGP neighbor state on the backup Routing Engine shows as idle. On switchover, the BGP session is reestablished from the new master Routing Engine.

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



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

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

Nonstop Active Routing Layer 2 Circuit and VPLS Support

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

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

Nonstop Active Routing PIM Support

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



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

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



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

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

Supported features:

- Auto-RP



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

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



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

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

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

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

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

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

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

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

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

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

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

Incompatible features: Nonstop active routing does not support the following features, and you cannot configure them on a router enabled for PIM nonstop active routing. The commit operation fails if the configuration includes both nonstop active routing and one or more of these features:

- Next-generation MVPNs with PIM provider tunnels

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

Nonstop Active Routing MSDP Support

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

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

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

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

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

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

Nonstop Active Routing Support for RSVP-TE LSPs

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

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

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

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

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

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

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

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

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

- State timeouts might take additional time during nonstop active routing switchover. For example, if a switchover occurs after a neighbor has missed sending two hello messages to the master, the new master Routing Engine waits for another three hello periods before timing out the neighbor.
- On the RSVP ingress router, if you configure auto-bandwidth functionality, the bandwidth adjustment timers are set in the new master after the switchover. This causes a one-time increase in the length of time required for the bandwidth adjustment after the switchover occurs.
- RSVP ingress LSPs that have BFD liveness detection enabled on them do not come up on the backup Routing Engine during the switchover. Such BFD-enabled LSPs have to be reestablished after the switchover.
- Backup LSPs —LSPs that are established between the point of local repair (PLR) and the merge point after a node or link failure—are not preserved during a Routing Engine switchover.
- When nonstop active routing is enabled, graceful restart is not supported. However, graceful restart helper mode is supported.

**Related
Documentation**

- [Nonstop Active Routing Concepts on page 15](#)
- [Configuring Nonstop Active Routing on page 137](#)

CHAPTER 4

Graceful Restart

- [Graceful Restart Concepts on page 29](#)
- [Graceful Restart System Requirements on page 30](#)
- [Aggregate and Static Routes on page 31](#)
- [Graceful Restart and Routing Protocols on page 31](#)
- [Graceful Restart and MPLS-Related Protocols on page 34](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 35](#)
- [Graceful Restart on Logical Systems on page 36](#)

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).
- 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 on page 30](#)
- [Aggregate and Static Routes on page 31](#)
- [Graceful Restart and Routing Protocols on page 31](#)
- [Graceful Restart and MPLS-Related Protocols on page 34](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 35](#)
- [Graceful Restart on Logical Systems on page 36](#)
- [Example: Configuring Graceful Restart on page 91](#)
- Configuring Graceful Restart for QFabric Systems

Graceful Restart System Requirements

Graceful restart is supported on all routing platforms. To implement graceful restart for particular features, your system must meet these minimum requirements:

- Junos OS Release 5.3 or later for aggregate route, BGP, IS-IS, OSPF, RIP, RIPng, or static route graceful restart.
- Junos OS Release 5.5 or later for RSVP on egress provider edge (PE) routers.
- Junos OS Release 5.5 or later for LDP graceful restart.
- Junos OS Release 5.6 or later for the CCC, TCC, Layer 2 VPN, or Layer 3 VPN implementations of graceful restart.
- Junos OS Release 6.1 or later for RSVP graceful restart on ingress PE routers.
- Junos OS Release 6.4 or later for PIM sparse mode graceful restart.
- Junos OS Release 7.4 or later for ES-IS graceful restart (J Series Services Routers).
- Junos OS Release 8.5 or later for BFD session (helper mode only)—If a node is undergoing a graceful restart and its BFD sessions are distributed to the Packet Forwarding Engine, the peer node can help the peer with the graceful restart.
- Junos OS Release 9.2 or later for BGP to support helper mode without requiring that graceful restart be configured.

Related Documentation

- [Graceful Restart Concepts on page 29](#)

Aggregate and Static Routes

When you include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level, any static routes or aggregated routes that have been configured are protected. Because no helper router assists in the restart, these routes are retained in the forwarding table while the router restarts (rather than being discarded or refreshed).

Related Documentation

- [Graceful Restart Concepts on page 29](#)
- [Graceful Restart System Requirements on page 30](#)
- [Enabling Graceful Restart on page 79](#)
- [Verifying Graceful Restart Operation on page 234](#)
- [Example: Configuring Graceful Restart on page 91](#)

Graceful Restart and Routing Protocols

This section covers the following topics:

- [BGP on page 31](#)
- [ES-IS on page 32](#)
- [IS-IS on page 32](#)
- [OSPF and OSPFv3 on page 32](#)
- [PIM Sparse Mode on page 33](#)
- [RIP and RIPng on page 33](#)

BGP

When a router enabled for BGP graceful restart restarts, it retains BGP peer routes in its forwarding table and marks them as stale. However, it continues to forward traffic to other peers (or receiving peers) during the restart. To reestablish sessions, the restarting router sets the "restart state" bit in the BGP OPEN message and sends it to all participating peers. The receiving peers reply to the restarting router with messages containing end-of-routing-table markers. When the restarting router or switch receives all replies from the receiving peers, the restarting router performs route selection, the forwarding table is updated, and the routes previously marked as stale are discarded. At this point, all BGP sessions are reestablished and the restarting peer can receive and process BGP messages as usual.

While the restarting router does its processing, the receiving peers also temporarily retain routing information. When a receiving peer detects a TCP transport reset, it retains the routes received and marks the routes as stale. After the session is reestablished with the restarting router or switch, the stale routes are replaced with updated route information.

ES-IS

When graceful restart for ES-IS is enabled, the routes to end systems or intermediate systems are not removed from the forwarding table. The adjacencies are reestablished after restart is complete.



NOTE: ES-IS is supported only on the J Series Services Router.

IS-IS

Normally, IS-IS routers move neighbor adjacencies to the down state when changes occur. However, a router enabled for IS-IS graceful restart sends out Hello messages with the Restart Request (RR) bit set in a restart type length value (TLV) message. This indicates to neighboring routers that a graceful restart is in progress and to leave the IS-IS adjacency intact. The neighboring routers must interpret and implement restart signaling themselves. Besides maintaining the adjacency, the neighbors send complete sequence number PDUs (CSNPs) to the restarting router and flood their entire database.

The restarting router never floods any of its own link-state PDUs (LSPs), including pseudonode LSPs, to IS-IS neighbors while undergoing graceful restart. This enables neighbors to reestablish their adjacencies without transitioning to the down state and enables the restarting router to reinitiate a smooth database synchronization.

OSPF and OSPFv3

When a router enabled for OSPF graceful restart restarts, it retains routes learned before the restart in its forwarding table. The router does not allow new OSPF link-state advertisements (LSAs) to update the routing table. This router continues to forward traffic to other OSPF neighbors (or helper routers), and sends only a limited number of LSAs during the restart period. To reestablish OSPF adjacencies with neighbors, the restarting router must send a grace LSA to all neighbors. In response, the helper routers enter helper mode and send an acknowledgement back to the restarting router. If there are no topology changes, the helper routers continue to advertise LSAs as if the restarting router had remained in continuous OSPF operation.

When the restarting router receives replies from all the helper routers, the restarting router selects routes, updates the forwarding table, and discards the old routes. At this point, full OSPF adjacencies are reestablished and the restarting router receives and processes OSPF LSAs as usual. When the helper routers no longer receive grace LSAs from the restarting router or the topology of the network changes, the helper routers also resume normal operation.



NOTE: For more information about the standard helper mode implementation, see RFC 3623, *Graceful OSPF Restart*.

Starting with Release 11.3, Junos OS supports the restart signaling-based helper mode for OSPF graceful restart configurations. The helper modes, both standard and restart

signaling-based, are enabled by default. In restart signaling-based helper mode implementations, the restarting router relays the restart status to its neighbors only after the restart is complete. When the restart is complete, the restarting router sends hello messages to its helper routers with the **restart signal (RS)** bit set in the hello packet header. When a helper router receives a hello packet with the **RS** bit set in the header, the helper router returns a hello message to the restarting router. The reply hello message from the helper router contains the **ResyncState** flag and the **ResyncTimeout** timer that enable the restarting router to keep track of the helper routers that are syncing up with it. When all helpers complete the synchronization, the restarting router exits the restart mode.



NOTE:

For more information about restart signaling-based graceful restart helper mode implementation, see RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*, RFC 4812, *OSPF Restart Signaling*, and RFC 4813, *OSPF Link-Local Signaling*.

Restart signaling-based graceful restart helper mode is not supported for OSPFv3 configurations.

PIM Sparse Mode

PIM sparse mode uses a mechanism called a *generation identifier* to indicate the need for graceful restart. Generation identifiers are included by default in PIM hello messages. An initial generation identifier is created by each PIM neighbor to establish device capabilities. When one of the PIM neighbors restarts, it sends a new generation identifier to its neighbors. All neighbors that support graceful restart and are connected by point-to-point links assist by sending multicast updates to the restarting neighbor.

The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires. If the neighbors do not support graceful restart or connect to each other using multipoint interfaces, the restarting router uses the restart interval timer to define the restart period.

RIP and RIPng

When a router enabled for RIP graceful restart restarts, routes that have been configured are protected. Because no helper router assists in the restart, these routes are retained in the forwarding table while the router restarts (rather than being discarded or refreshed).

Related Documentation

- [Graceful Restart Concepts on page 29](#)
- [Graceful Restart System Requirements on page 30](#)
- [Configuring Routing Protocols Graceful Restart on page 80](#)
- [Verifying Graceful Restart Operation on page 234](#)
- [Example: Configuring Graceful Restart on page 91](#)

Graceful Restart and MPLS-Related Protocols

This section contains the following topics:

- [LDP on page 34](#)
- [RSVP on page 34](#)
- [CCC and TCC on page 35](#)

LDP

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

The reconnect time is configured in Junos OS as 60 seconds and is not user-configurable. The reconnect time is how long the helper router waits for the restarting router to establish a connection. If the connection is not established within the reconnect interval, graceful restart for the LDP session is terminated. The maximum reconnect time is 120 seconds and is not user-configurable. The maximum reconnect time is the maximum value that a helper router accepts from its restarting neighbor.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, so it can continue to forward traffic.

You can configure LDP graceful restart both in the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and for a specific routing instance only.

RSVP

RSVP graceful restart enables a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

CCC and TCC

CCC and TCC graceful restart enables Layer 2 connections between customer edge (CE) routers to restart gracefully. These Layer 2 connections are configured with the **remote-interface-switch** or **lsp-switch** statements. Because these CCC and TCC connections have an implicit dependency on RSVP LSPs, graceful restart for CCC and TCC uses the RSVP graceful restart capabilities.

RSVP graceful restart must be enabled on the provider edge (PE) routers and provider (P) routers to enable graceful restart for CCC and TCC. Also, because RSVP is used as the signaling protocol for signaling label information, the neighboring router must use helper mode to assist with the RSVP restart procedures.

Related Documentation

- [Graceful Restart Concepts on page 29](#)
- [Graceful Restart System Requirements on page 30](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 86](#)
- [Example: Configuring Graceful Restart on page 91](#)

Graceful Restart and Layer 2 and Layer 3 VPNs

VPN graceful restart uses three types of restart functionality:

1. BGP graceful restart functionality is used on all PE-to-PE BGP sessions. This affects sessions carrying any service signaling data for network layer reachability information (NLRI), for example, an IPv4 VPN or Layer 2 VPN NLRI.
2. OSPF, IS-IS, LDP, or RSVP graceful restart functionality is used in all core routers. Routes added by these protocols are used to resolve Layer 2 and Layer 3 VPN NLRI.
3. Protocol restart functionality is used for any Layer 3 protocol (RIP, OSPF, LDP, and so on) used between the PE and customer edge (CE) routers. This does not apply to Layer 2 VPNs because Layer 2 protocols used between the CE and PE routers do not have graceful restart capabilities.

Before VPN graceful restart can work properly, all of the components must restart gracefully. In other words, the routers must preserve their forwarding states and request neighbors to continue forwarding to the router in case of a restart. If all of the conditions are satisfied, VPN graceful restart imposes the following rules on a restarting router:

- The router must wait to receive all BGP NLRI information from other PE routers before advertising routes to the CE routers.
- The router must wait for all protocols in all routing instances to converge (or complete the restart process) before it sends CE router information to other PE routers. In other words, the router must wait for all instance information (whether derived from local configuration or advertisements received from a remote peer) to be processed before it sends this information to other PE routers.

- The router must preserve all forwarding state in the **instance.mpls.0** tables until the new labels and transit routes are allocated and announced to other PE routers (and CE routers in a carrier-of-carriers scenario).

If any condition is not met, VPN graceful restart does not succeed in providing uninterrupted forwarding between CE routers across the VPN infrastructure.

**Related
Documentation**

- [Graceful Restart Concepts on page 29](#)
- [Graceful Restart System Requirements on page 30](#)
- [Configuring Logical System Graceful Restart on page 89](#)
- [Verifying Graceful Restart Operation on page 234](#)
- [Example: Configuring Graceful Restart on page 91](#)

Graceful Restart on Logical Systems

Graceful restart for a logical system functions much as graceful restart does in the main router. The only difference is the location of the **graceful-restart** statement:

- For a logical system, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-options]** hierarchy level.
- For a routing instance inside a logical system, include the **graceful-restart** statement at both the **[edit logical-systems *logical-system-name* routing-options]** and **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]** hierarchy levels.

**Related
Documentation**

- [Graceful Restart Concepts on page 29](#)
- [Graceful Restart System Requirements on page 30](#)
- [Configuring Logical System Graceful Restart on page 89](#)
- [Verifying Graceful Restart Operation on page 234](#)
- [Example: Configuring Graceful Restart on page 91](#)

CHAPTER 5

Unified ISSU

- [Upgrading Routers Using ISSU on page 37](#)
- [Unified ISSU Concepts on page 37](#)
- [Unified ISSU System Requirements on page 42](#)

Upgrading Routers Using ISSU

Unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. ISSU is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

For additional information about using ISSU, see the [Junos OS High Availability Guide](#).

Unified ISSU Concepts

A unified in-service software upgrade (unified ISSU) enables you to upgrade between two different Junos OS Releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is only supported on dual Routing Engine platforms. In addition, the graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

A unified ISSU provides the following benefits:

- Eliminates network downtime during software image upgrades
- Reduces operating costs, while delivering higher service levels
- Allows fast implementation of new features



NOTE: The master Routing Engine and backup Routing Engine must be running the same software version before you can perform a unified ISSU.

You cannot take any PICs online or offline during a unified ISSU.



NOTE: You can verify the unified ISSU-compatibility of the software, hardware, and the configuration on a device by issuing the `request system software validate in-service-upgrade` command. This command runs the validation checks, and shows whether the operating system, device components, and configurations are ISSU compatible or not. For more information about the `request system software validate in-service-upgrade` command, see Junos OS Operational Mode Commands.



NOTE: Unicast RPF-related statistics are not saved across a unified ISSU, and the unicast RPF counters are reset to zero during a unified ISSU.

To perform a unified ISSU, complete the following steps:

1. Enable graceful Routing Engine switchover and nonstop active routing. Verify that the Routing Engines and protocols are synchronized.
2. Download the new software package from the Juniper Networks Support website and then copy the package to the router.
3. Issue the `request system software in-service-upgrade` command on the master Routing Engine.

A Junos OS Release package comprises three distinct systems:

- Juniper Networks Operating System, which provides system control and all the features and functions of the Juniper Networks router that executes in the Routing Engines
- Juniper Networks Packet Forwarding Engine, which supports the high-performance traffic forwarding and packet handling capabilities
- Interface control

After the `request system software in-service-upgrade` command is issued, the following process occurs.

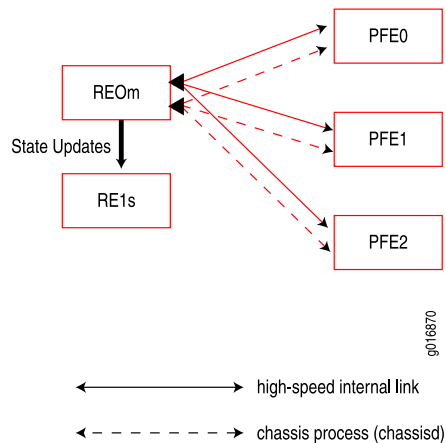


NOTE: In the illustrations, a solid line indicates the high-speed internal link between a Routing Engine and a Packet Forwarding Engine. A dotted line indicates the chassis process (`chassisd`), another method of communication between a Routing Engine and a Packet Forwarding Engine. RE0m and RE1s indicate master and backup (or standby) Routing Engines, respectively.



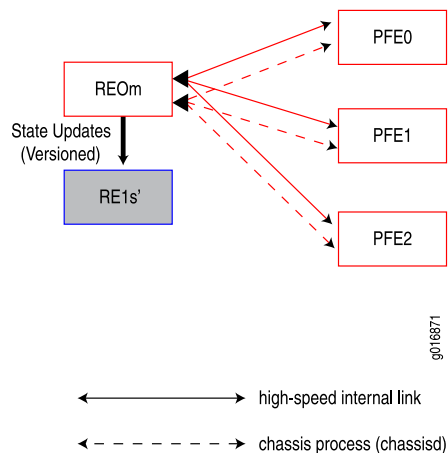
NOTE: The following process pertains to all supported routing platforms except the TX Matrix router. For information about the unified ISSU process on the TX Matrix router, see *Unified ISSU Process on the TX Matrix Router*. On M320 and T320 routers and on T640 and T1600 routers, the Packet Forwarding Engine resides on an FPC. However, on an M120 router, the Forwarding Engine Board (FEB) replaces the functions of a Packet Forwarding Engine. In the illustrations and steps, when considering an M120 router, you can regard the PFE as an FPC. As an additional step on an M120 router, after the FPCs and PICs have been upgraded, the FEBs are upgraded.

1. The master Routing Engine validates the router configuration to ensure that it can be committed using the new software version. Checks are made for disk space available for the `/var` file system on both Routing Engines, unsupported configurations, and for unsupported Physical Interface Cards (PICs). If there is not sufficient disk space available on either of the Routing Engines, the unified ISSU process fails and returns an error message saying that the Routing Engine does not have enough disk space available. However, unsupported PICs do not prevent a unified ISSU. The software issues a warning to indicate that these PICs will restart during the upgrade. Similarly, an unsupported protocol configuration does not prevent a unified ISSU. The software issues a warning that packet loss may occur for the protocol during the upgrade.

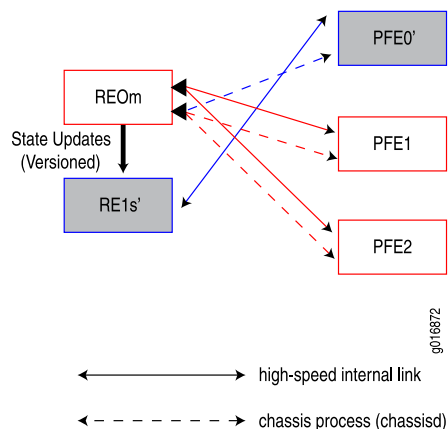


2. When the validation succeeds, the kernel state synchronization daemon (ksyncd) synchronizes the kernel on the backup Routing Engine with the master Routing Engine.
3. The backup Routing Engine is upgraded with the new software image. Before being upgraded, the backup Routing Engine gets the configuration file from the master Routing Engine and validates the configuration to ensure that it can be committed using the new software version. After being upgraded, it is resynchronized with the

master Routing Engine. In the illustration, an apostrophe (') indicates the device is running the new version of software.




4. The chassis process (chassisd) on the master Routing Engine prepares other software processes for the unified ISSU. When all the processes are ready, chassisd sends an ISSU_PREPARE message to the Flexible PIC Concentrators (FPCs) installed in the router.
5. The Packet Forwarding Engine on each FPC saves its state and downloads the new software image from the backup Routing Engine. Next, each Packet Forwarding Engine sends an ISSU_READY message to the chassis process (chassisd).




6. After receiving an ISSU_READY message from a Packet Forwarding Engine, the chassis process (chassisd) sends an ISSU_REBOOT message to the FPC on which the Packet Forwarding Engine resides. The FPC reboots with the new software image. After the FPC is rebooted, the Packet Forwarding Engine restores the FPC state and a high-speed internal link is established with the backup Routing Engine running the new software. The chassis process (chassisd) is also reestablished with the master Routing Engine.
7. After all Packet Forwarding Engines have sent a READY message using the chassis process (chassisd) on the master Routing Engine, other software processes are

prepared for a Routing Engine switchover. The system is ready for a switchover at this point.

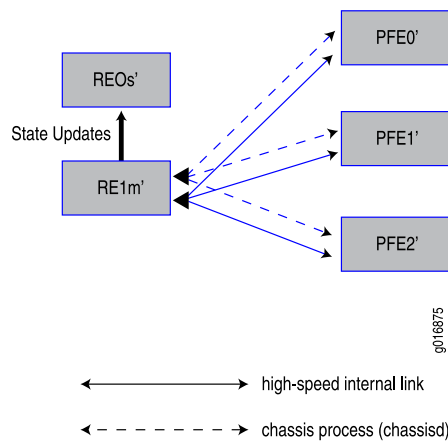


NOTE: In the case of an M120 router, the FEBs are upgraded at this point. When all FEBs have been upgraded, the system is ready for a switchover.

8. The Routing Engine switchover occurs and the backup Routing Engine becomes the new master Routing Engine.



9. The new backup Routing Engine is now upgraded to the new software image. (This step is skipped if the **no-old-master-upgrade** option is specified.)



10. When the backup Routing Engine has been successfully upgraded, the unified ISSU is complete.

Related Documentation

- [Unified ISSU Process on the TX Matrix Router](#)
- [Unified ISSU System Requirements on page 42](#)
- [Best Practices on page 153](#)
- [Before You Begin on page 154](#)
- [Performing a Unified ISSU on page 157](#)

Unified ISSU System Requirements

This section contains the following topics:

- [Unified ISSU Junos OS Release Support on page 42](#)
- [Unified ISSU Platform Support on page 43](#)
- [Unified ISSU Protocol Support on page 43](#)
- [Unified ISSU Feature Support on page 45](#)
- [Unified ISSU PIC Support on page 45](#)
- [Unified ISSU Support on MX Series 3D Universal Edge Routers on page 53](#)

Unified ISSU Junos OS Release Support

In order to perform a unified ISSU, your router must be running a Junos OS Release that supports unified ISSU for the specific platform. See [“Unified ISSU Platform Support” on page 43](#). You can use unified ISSU to upgrade from an ISSU-capable software release to a newer software release. However, note that:

- The unified ISSU process is aborted if the Junos OS version specified for installation is a version earlier than the one currently running on the device. To downgrade from an

ISSU-capable release to a previous software release (ISSU-capable or not), use the **request system add** command. Unlike an upgrade using the unified ISSU process, a downgrade using the **request system add** command can cause network disruptions and loss of data. For more information about the use of the **request system add** command, see the Installation and Upgrade Guide.

- The unified ISSU process is aborted if the specified upgrade has conflicts with the current configuration, components supported, and so forth.
- Unified ISSU does not support extension application packages developed using the Junos SDK.

Unified ISSU Platform Support

Table 7 on page 43 lists the platforms on which a unified ISSU is supported.

Table 7: Unified ISSU Platform Support

Routing Platform	Junos OS Release
M120 router	9.2 or later
M320 router	9.0 or later
M10i router with Enhanced Compact Forwarding Engine Board (CFEB-E)	9.5 or later
MX Series devices	9.3 or later
NOTE: Unified ISSU for MX Series routers does not support IEEE 802.1ag OAM and IEEE 802.3ah protocols.	11.2 or later on MX Series 3D Universal Edge Routers (with Trio Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces).
T320 router	9.0 or later
T640 router	9.0 or later
T1600 router	9.1 or later
TX Matrix router	9.3 or later

Unified ISSU Protocol Support

Unified ISSU is dependent on nonstop active routing. Table 8 on page 43 lists the protocols that are supported during a unified ISSU.

Table 8: Unified ISSU Protocol Support

Protocol	Junos OS Release
BGP	9.0 or later

Table 8: Unified ISSU Protocol Support (*continued*)

Protocol	Junos OS Release
DHCP access model (subscriber access)	11.2 or later
IS-IS	9.0 or later
LDP	9.0 or later
LDP-based virtual private LAN service (VPLS)	9.3 or later
Layer 2 circuits	9.2 or later
Layer 3 VPNs using LDP	9.2 or later
Link Aggregation Control Protocol (LACP) on MX Series routers	9.4 or later
OSPF/OSPFv3	9.0 or later
PPPoE access model (subscriber access)	11.4 or later
Protocol Independent Multicast (PIM)	9.3 or later
Routing Information Protocol (RIP)	9.1 or later

Unified ISSU Support for the Layer 2 Control Protocol Process

Unified ISSU supports the Layer 2 Control Protocol process (l2cpd) on MX Series 3D Universal Edge Routers. In a Layer 2 bridge environment, spanning tree protocols share information about port roles, bridge IDs, and root path costs between bridges using special data frames called Bridge Protocol Data Units (BPDUs). The transmission of BPDUs is controlled by the l2cpd process. Transmission of hello BPDUs is important in maintaining adjacencies on the peer systems.

The transmission of periodic packets on behalf of the l2cpd process is carried out by periodic packet management (PPM), which, by default, is configured to run on the Packet Forwarding Engine. The ppm process on the Packet Forwarding Engine ensures that the BPDUs are transmitted even when the l2cpd process control plane is unavailable, and keeps the remote adjacencies alive during unified ISSU. However, if you want the distributed PPM (ppmd) process to run on the Routing Engine instead of the Packet Forwarding Engine, you can disable the ppm process on the Packet Forwarding Engine, by including the **no-delegate-processing** statement at the [edit routing-options ppm] hierarchy level.



NOTE: The `delegate-processing` statement at the `[edit routing-options ppm]` hierarchy level, which was used to enable the `ppmd` process on the Packet Forwarding Engine in Junos OS Release 9.3 and earlier, has been deprecated as the `ppmd` process is enabled on the Packet Forwarding Engine by default in Junos OS Release 9.4 and later.

Unified ISSU enhancements and nonstop active bridging support for the `l2cpd` process ensure that the new master Routing Engine is able to take control during unified ISSU without any disruptions in the control plane and minimize the disruptions in the Layer 2 data plane during unified ISSU.

Unified ISSU Feature Support

Unified ISSU supports most Junos OS features in Junos OS Release 9.0. However, the following constraints apply:

- Link Aggregation Control Protocol (LACP)—Link changes are not processed until after the unified ISSU is complete.
- Automatic Protection Switching (APS)—Network changes are not processed until after the unified ISSU is complete.
- Ethernet Operation, Administration, and Management (OAM) as defined by IEEE 802.3ah and by IEEE 802.1ag—When a Routing Engine switchover occurs, the OAM hello times out, triggering protocol convergence.
- Ethernet circuit cross-connect (CCC) encapsulation—Circuit changes are not processed until after the unified ISSU is complete.
- Logical systems—On routers that have logical systems configured on them, only the master logical system supports unified ISSU.

Unified ISSU PIC Support

The following sections list the Physical Interface Cards (PICs) that are supported during a unified ISSU.

- [PIC Considerations on page 46](#)
- [SONET/SDH PICs on page 46](#)
- [Fast Ethernet and Gigabit Ethernet PICs on page 48](#)
- [Channelized PICs on page 49](#)
- [Tunnel Services PICs on page 50](#)
- [ATM PICs on page 50](#)
- [Serial PICs on page 51](#)
- [DS3, E1, E3, and T1 PICs on page 51](#)
- [Enhanced IQ PICs on page 52](#)
- [Enhanced IQ2 Ethernet Services Engine \(ESE\) PIC on page 52](#)



NOTE: 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.

PIC Considerations

Take the following PIC restrictions into consideration before performing a unified ISSU:

- **Unsupported PICs**—If a PIC is not supported by unified ISSU, at the beginning of the upgrade the software issues a warning that the PIC will be brought offline. After the PIC is brought offline and the ISSU is complete, the PIC is brought back online with the new firmware.
- **PIC combinations**—For some PICs, newer Junos OS services can require significant Internet Processor ASIC memory, and some configuration rules might limit certain combinations of PICs on particular platforms. With a unified ISSU:
 - If a PIC combination is not supported by the software version that the router is being upgraded from, the upgrade will be aborted.
 - If a PIC combination is not supported by the software version to which the router is being upgraded, the in-service software upgrade will abort, even if the PIC combination is supported by the software version from which the router is being upgraded.
- **Interface statistics**—Interface statistics might be incorrect because:
 - During bootup of the new microkernel on the Packet Forwarding Engine (PFE), host-bound traffic is not handled and might be dropped, causing packet loss.
 - During the hardware update of the Packet Forwarding Engine and its interfaces, traffic is halted and discarded. (The duration of the hardware update depends on the number and type of interfaces and on the router configuration.)
 - During a unified ISSU, periodic statistics collection is halted. If hardware counters saturate or wrap around, the software does not display accurate interface statistics.
- **CIR oversubscription**—If oversubscription of committed rate information (CIR) is configured on logical interfaces:
 - And the sum of the CIR exceeds the physical interface's bandwidth, after a unified in-service software upgrade is performed, each logical interface might not be given its original CIR.
 - And the sum of the delay buffer rate configured on logical interfaces exceeds the physical interface's bandwidth, after a unified in-service software upgrade is performed, each logical interface might not receive its original delay-buffer-rate calculation.

SONET/SDH PICs

Table 9 on page 47 lists the SONET/SDH PICs that are supported during a unified ISSU.

Table 9: Unified ISSU PIC Support: SONET/SDH

PIC Type	Number of Ports	Model Number	Router
OC3c/STM1	4-port	PB-4OC3-SON-MM—(EOL)	M120 M320, T320, T640, T1600
		PB-4OC3-SON-SMIR—(EOL)	
		PE-4OC3-SON-MM—(EOL)	M10i
		PE-4OC3-SON-SMIR—(EOL)	
	2-port	PE-2OC3-SON-MM—(EOL)	
		PE-2OC3-SON-SMIR—(EOL)	
OC3c/STM1 with SFP	2-port	PE-2OC3-SON-SFP	M10i
OC3c/STM1, SFP (Multi-Rate)	4 OC3 ports, 4 OC12 ports	PB-4OC3-4OC12-SON-SFP	M120 M320, MX Series, T320, T640, T1600
	4 OC3 ports, 1 OC12 port	PB-4OC3-1OC12-SON-SFP	
		PB-4OC3-1OC12-SON2-SFP	M10i
		PE-4OC3-1OC12-SON-SFP	
OC12c/STM4	1-port	PE-1OC12-SON-SFP	M10i
		PE-1OC12-SON-MM—(EOL)	
		PE-1OC12-SON-SMIR—(EOL)	
		PB-1OC12-SON-MM—(EOL)	
		PB-1OC12-SON-SMIR—(EOL)	M120, M320, T320, T640, T1600, TX Matrix
	4-port	PB-4OC12-SON-MM	
		PB-4OC12-SON-SMIR	
OC12c/STM4, SFP	1-port	PB-1OC12-SON-SFP	M120, M320, T320, T640, T1600, TX Matrix
OC48c/STM16, SFP	1-port	PB-1OC48-SON-SFP	M120, M320, MX Series, T320, T640, T1600, TX Matrix
		PB-1OC48-SON-B-SFP	
	4-port	PC-4OC48-SON-SFP	
OC192/STM64	1-port	PC-1OC192-SON-VSR	MX Series routers

Table 9: Unified ISSU PIC Support: SONET/SDH (*continued*)

PIC Type	Number of Ports	Model Number	Router
OC192/STM64, XFP	1-port	PC-1OC192-SON-LR	M320, T320, T640, T1600
		PC-1OC192-SON-SR2	
		PC-1OC192-VSR	
OC192/STM64, XFP	4-port	PD-4OC192-SON-XFP	M120, T640, T1600
	1-port	PC-1OC192-SON-XFP	MX Series routers
OC768/STM256	1-port	PD-1OC768-SON-SR	T640, T1600

Fast Ethernet and Gigabit Ethernet PICs

Table 10 on page 48 lists the Fast Ethernet and Gigabit Ethernet PICs that are supported during a unified ISSU.



NOTE: Starting with Junos OS Release 9.2, new Ethernet IQ2 PIC features might cause the software to reboot the PIC when a unified ISSU is performed. For information about applicable new Ethernet IQ2 PIC features, refer to the release notes for the specific Junos OS Release.

Table 10: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet

PIC Type	Number of Ports	Model Number	Router
Fast Ethernet	4	PB-4FE-TX	M120, M320, T320, T640, T1600, TX Matrix
		PE-4FE-TX	M10i
	8	PB-8FE-FX	M120, M320
		PE-8FE-FX	M10i
	12	PB-12FE-TX-MDI	M120, M320, T320
		PB-12FE-TX-MDIX	
		PE-12FE-TX-MDI	M10i
		PE-12FE-TX-MDIX	
	48	PB-48FE-TX-MDI	M120, M320, T320
		PB-48FE-TX-MDIX	

Table 10: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet (*continued*)

PIC Type	Number of Ports	Model Number	Router
Gigabit Ethernet, SFP	1	PE-1GE-SFP	M10i
		PB-1GE-SFP	M120, M320, T320, T640, T1600, TX Matrix
	2	PB-2GE-SFP	
	4	PB-4GE-SFP	
	10	PC-10GE-SFP	
Gigabit Ethernet IQ, SFP	1	PE-1GE-SFP-QPP	M10i
		PB-1GE-SFP-QPP	M120, M320, T320, T640, T1600, TX Matrix
	2	PB-2GE-SFP-QPP	
Gigabit Ethernet IQ2, SFP	4	PB-4GE-TYPE1-SFP-IQ2	M120, M320, T320, T640, T1600, TX Matrix
	8	PB-8GE-TYPE2-SFP-IQ2	
		PC-8GE-TYPE3-SFP-IQ2	
Gigabit Ethernet IQ2, XFP	1	PC-1XGE-TYPE3-XFP-IQ2	M120, M320, T320, T640, T1600, TX Matrix
10-Gigabit Ethernet, XENPAK	1	PC-1XGE-XENPAK	M120, M320, T320, T640, T1600, TX Matrix
10-Gigabit Ethernet, DWDM	1	PC-1XGE-DWDM-CBAND	M120, M320, T320, T640, T1600, TX Matrix
10-Gigabit Ethernet	4	PD-4XGE-XFP <i>NOTE:</i> This PIC goes offline during a unified ISSU if the PIC is inserted on T-1600-FPC4-ES with revision number less than 13.	T640, T1600, TX Matrix, TX Matrix Plus
	10	PD-5-10XGE-SFPP	T640, T1600

Channelized PICs

Table 11 on page 50 lists the channelized PICs that are supported during a unified ISSU.

Table 11: Unified ISSU PIC Support: Channelized

PIC Type	Number of Ports	Model Number	Platform
Channelized E1 IQ	10	PB-10CHE1-RJ48-QPP	M120, M320, T320, T640, T1600, TX Matrix
		PE-10CHE1-RJ48-QPP-N	M10i
Channelized T1 IQ	10	PB-10CHT1-RJ48-QPP	M320, T320, T640, T1600
		PE-10CHT1-RJ48-QPP	M10i
Channelized OC IQ	1	PB-1CHOC12SMIR-QPP	M120, M320, T320, T640, T1600, TX Matrix
		PB-1CHSTM1-SMIR-QPP	
		PB-1CHOC3-SMIR-QPP	
		PE-1CHOC12SMIR-QPP	M10i
		PE-1CHOC3-SMIR-QPP	
Channelized DS3 to DS0 IQ	4	PB-4CHDS3-QPP	M120, M320, T320, T640, T1600, TX Matrix
		PE-4CHDS3-QPP	M10i
Channelized STM 1	1	PE-1CHSTM1-SMIR-QPP	M10i

Tunnel Services PICs

Table 12 on page 50 lists the Tunnel Services PICs that are supported during a unified ISSU.

Table 12: Unified ISSU PIC Support: Tunnel Services

PIC Type	Model Number	Platform
1-Gbps Tunnel	PE-TUNNEL	M10i
	PB-TUNNEL-1	M120, M320, T320, T640, T1600, TX Matrix
4-Gbps Tunnel	PB-TUNNEL	
10-Gbps Tunnel	PC-TUNNEL	

ATM PICs

Table 13 on page 51 lists the ATM PICs that are supported during a unified ISSU. This includes support on Enhanced III FPCs.

Table 13: Unified ISSU PIC Support: ATM

PIC Type	Number of Ports	Model Number	Platform
DS3	4	PB-4DS3-ATM2	M120, M320, T320, T640, T1600, TX Matrix
		PE-4DS3-ATM2	M10i
E3	4	PB-4E3-ATM2	M120, M320, T320, T640, T1600, TX Matrix
	2	PE-2E3-ATM2	M10i
OC3/STM1	2	PB-2OC3-ATM2-MM	M120, M320, T320, T640, T1600, TX Matrix
		PB-2OC3-ATM2-SMIR	
		PE-2OC3-ATM2-MM	M10i
		PE-2OC3-ATM2-SMIR	
OC12/STM4	1	PB-1OC12-ATM2-MM	M120, M320, T320, T640, T1600, TX Matrix
		PB-1OC12-ATM2-SMIR	
	2	PB-2OC12-ATM2-MM	M120, M320, T320, T640, T1600, TX Matrix
		PB-2OC12-ATM2-SMIR	
	1	PE-1OC12-ATM2-MM	M10i
		PE-1OC12-ATM2-SMIR	
OC48/STM16	1	PB-1OC48-ATM2-SFP	M120, M320, T320, T640, T1600, TX Matrix

Serial PICs

Unified ISSU supports the following 2-port EIA-530 serial PICs:

- PB-2EIA530 on M320 routers with Enhanced III FPCs, and on M120 routers
- PE-2EIA530 on M10i routers

DS3, E1, E3, and T1 PICs

Unified ISSU supports the following PICs on M120, M320, and T320 routers; T640 and T1600 routers; and the TX Matrix router:

- 4-Port DS3 PIC (PB-4DS3)
- 4-Port E1 Coaxial PIC (PB-4E1-COAX)
- 4-Port E1 RJ48 PIC (PB-4E1-RJ48)

- 4-port E3 IQ PIC (PB-4E3-QPP)
- 4-Port T1 PIC (PB-4T1-RJ48)

Unified ISSU supports the following PICs on M10i routers:

- 2-Port DS3 PIC (PE-2DS3)
- 4-Port DS3 PIC (PE-4DS3)
- 4-Port E1 PICs (PE-4E1-COAX and PE-4E1-RJ48)
- 2-Port E3 PIC (PE-2E3)
- 4-Port T1 PIC (PE-4T1-RJ48)
- 4-Port E3 IQ PIC (PE-4E3-QPP)

Enhanced IQ PICs

Unified ISSU supports the following PICs on M120, M320, and T320 routers; T640 and T1600 routers; and the TX Matrix router:

- 1-port Channelized OC12/STM4 enhanced IQ PIC (PB-1CHOC12-STM4-IQE-SFP)
- 1-port nonchannelized OC12/STM4 enhanced IQ PIC (PB-1OC12-STM4-IQE-SFP)
- 4-port Channelized DS3/E3 enhanced IQ PIC (PB-4CHDS3-E3-IQE-BNC)
- 4-port nonchannelized DS3/E3 enhanced IQ PIC (PB-4DS3-E3-IQE-BNC)
- 4-port nonchannelized SONET/SDH OC48/STM16 Enhanced IQ (IQE) PIC with SFP (PC-4OC48-STM16-IQE-SFP)

Unified ISSU supports 1-port Channelized OC48/STM16 Enhanced IQ (IQE) PIC with SFP (PB-1CHOC48-STM16-IQE-SFP) on MX Series routers.

Enhanced IQ2 Ethernet Services Engine (ESE) PIC

Unified ISSU supports the enhanced IQ2 ESE PICs listed in [Table 14 on page 52](#).

Table 14: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC

Model Number	Number of Ports	Platform
PC-8GE-TYPE3-SFP-IQ2E	8	M120, M320, T320, T640, and TX Matrix.
PB-8GE-TYPE2-SFP-IQ2E	8	M120, M320, T320, T640, and TX Matrix.
PB-4GE-TYPE1-SFP-IQ2E	4	M120, M320, T320, and T640.
PC-1XGE-TYPE3-XFP-IQ2E	1	M120, M320, T320, T640, and TX Matrix.
PB-1CHOC48-STM16-IQE	1	M120, M320, T320, T640, and TX Matrix.
PE-4GE-TYPE1-SFP-IQ2E	4	M10i.

Table 14: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC (*continued*)

Model Number	Number of Ports	Platform
PE-4GE-TYPE1-SFP-IQ2	4	M10i.

Unified ISSU Support on MX Series 3D Universal Edge Routers

The following sections list the Dense Port Concentrators (DPCs), Flexible PIC Concentrators (FPCs), Modular Port Concentrators (MPCs), and Modular Interface Cards (MICs) that are supported during a unified ISSU on MX Series 3D Universal Edge Routers.

- [Unified ISSU DPC and FPC Support on MX Series 3D Universal Edge Routers on page 53](#)
- [Unified ISSU MIC and MPC Support on MX Series 3D Universal Edge Routers on page 53](#)
- [Unified ISSU Limitation on MX Series 3D Universal Edge Routers on page 54](#)

Unified ISSU DPC and FPC Support on MX Series 3D Universal Edge Routers

Unified ISSU supports all Dense Port Concentrators (DPCs) except the Multiservices DPC on the MX Series routers. However, unified ISSU does not support either of the FPCs (FPC type 2, **MX-FPC2**, or FPC type 3, **MX-FPC3**) on the MX Series routers. For more information about DPCs and FPCs on MX Series routers, go to http://www.juniper.net/techpubs/en_US/release-independent/junos/information-products/pathway-pages/mx-series/.

Unified ISSU MIC and MPC Support on MX Series 3D Universal Edge Routers

Unified ISSU supports all the Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) listed in [Table 15 on page 53](#) and [Table 16 on page 54](#). Unified ISSU is not supported on MX80 routers nor in an MX Series Virtual Chassis.

In the MPCs on MX Series routers, interface-specific and firewall filter statistics are preserved across a unified ISSU. During the unified ISSU, counter and policer operations are disabled.

To preserve statistics across a unified ISSU on MX Series routers with MPC/MIC interfaces, the router stores the statistics data as binary large objects. The router collects the statistics before the unified ISSU is initialized, and restores the statistics after the unified ISSU completes. No statistics are collected during the unified ISSU process.

To verify that statistics are preserved across the unified ISSU, you can issue CLI operational commands such as **show interfaces statistics** after the unified ISSU completes.

Table 15: Unified ISSU Support: MX Series 3D Universal Edge Routers

MPC Type	Number of Ports	Model Number	Platform
30-Gigabit Ethernet MPC	—	MX-MPC1-3D	MX Series 3D Universal Edge Routers
30-Gigabit Ethernet Queuing MPC	—	MX-MPC1-3D-Q	MX Series 3D Universal Edge Routers

Table 15: Unified ISSU Support: MX Series 3D Universal Edge Routers (*continued*)

MPC Type	Number of Ports	Model Number	Platform
60-Gigabit Ethernet MPC	—	MX-MPC2-3D	MX Series 3D Universal Edge Routers
60-Gigabit Ethernet Queuing MPC	—	MX-MPC2-3D-Q	MX Series 3D Universal Edge Routers
60-Gigabit Ethernet Enhanced Queuing MPC	—	MX-MPC2-3D-EQ	MX Series 3D Universal Edge Routers
10-Gigabit Ethernet MPC with SFP+	16	MPC-3D-16XGE-SFPP	MX Series 3D Universal Edge Routers

Table 16: Unified ISSU Support: MX Series 3D Universal Edge Routers

MIC Type	Number of Ports	Model Number	Platform
Gigabit Ethernet MIC with SFP	20	MIC-3D-20GE-SFP	MX Series 3D Universal Edge Routers
10-Gigabit Ethernet MICs with XFP	2	MIC-3D-2XGE-XFP	MX Series 3D Universal Edge Routers
10-Gigabit Ethernet MICs with XFP	4	MIC-3D-4XGE-XFP	MX Series 3D Universal Edge Routers
Tri-Rate Copper Ethernet MIC	40	MIC-3D-40GE-TX	MX Series 3D Universal Edge Routers



NOTE: Note that unified ISSU is supported only by the MICs listed in [Table 16 on page 54](#).

Unified ISSU Limitation on MX Series 3D Universal Edge Routers

Unified in-service software upgrade (unified ISSU) is currently not supported when clock synchronization is configured for Synchronous Ethernet, Precision Time Protocol (PTP), and hybrid mode on MX80 3D Universal Edge Routers and on the MICs and MPCEs on MX240, MX480, and MX960 routers.

Related Documentation

- [Unified ISSU Concepts on page 37](#)
- [Unified ISSU Process on the TX Matrix Router](#)
- [Before You Begin on page 154](#)
- [Performing a Unified ISSU on page 157](#)

CHAPTER 6

VRRP

- [Understanding VRRP on page 55](#)
- [Junos OS Support for VRRPv3 on page 56](#)
- [Improving the Convergence Time for VRRP on page 59](#)

Understanding VRRP

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

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




NOTE: To minimize network traffic, VRRP is designed in such a way that only the router that is acting as the master sends out VRRP advertisements at any given point in time. The backup routers do not send any advertisement until and unless they take over mastership.

VRRP for IPv6 provides a much faster switchover to an alternate default router than IPv6 Neighbor Discovery (ND) procedures. Typical deployments use only one backup router.

Figure 7 on page 56 illustrates a basic VRRP topology. In this example, Routers A, B, and C are running VRRP and together they make up a virtual router. The IP address of this virtual router is 10.10.0.1 (the same address as the physical interface of Router A).

Figure 7: Basic VRRP



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

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*. VRRP for IPv6 is defined in Internet draft draft-ietf-vrrp-ipv6-spec-08.txt, *Virtual Router Redundancy Protocol for IPv6*. See also Internet draft draft-ietf-vrrp-unified-mib-06.txt, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6*.



NOTE: Even though VRRP, as defined in RFC 3768, does not support authentication, the Junos OS implementation of VRRP supports authentication as defined in RFC 2338. This support is achieved through the backward compatibility options in RFC 3768.

Related Documentation

- Understanding High Availability Features on Juniper Networks Routers
- [Configuring Basic VRRP Support on page 180](#)

Junos OS Support for VRRPv3

Prior to Junos OS Release 12.2, Junos OS supported RFC 3768, *Virtual Router Redundancy Protocol (VRRP)* and Internet draft draft-ietf-vrrp-ipv6-spec-08, *Virtual Router Redundancy Protocol for IPv6*. Starting with Junos OS Release 12.2, Junos OS supports Virtual Router Redundancy Protocol version 3 (VRRPv3). The support for VRRPv3 is implemented in compliance with RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6*. Junos OS Release 12.2 also supports VRRP MIB for VRRPv3.

The support for VRRP MIB for VRRPv3 is implemented in compliance with RFC 6527, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol Version 3 (VRRPv3)*.

When you configure VRRP for IPv4 or IPv6 networks, you can enable VRRPv3 by configuring the **version-3** statement at the **[edit protocols vrrp]** hierarchy level.



NOTE: When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with the previous versions of VRRP.

Understanding VRRPv3 Behavioral Differences

You must consider the following aspects when enabling VRRPv3 for your IPv4 or IPv6 networks:

- When VRRP for IPv6 is configured without enabling VRRPv3, the VRRP checksum is calculated according to section 5.3.8 of RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*. However, when VRRPv3 is enabled, the VRRP checksum is calculated according to section 5.3.7 of draft-ietf-vrrp-ipv6-spec-08.txt, *Virtual Router Redundancy Protocol for IPv6*. Therefore, when IPv6 VRRP packets are received or transmitted, the VRRP checksum is calculated according to:
 - RFC 3768, when VRRPv3 is *not* enabled.
 - draft-ietf-vrrp-ipv6-spec-08.txt, when VRRPv3 is enabled.
- The **tcpdump** utility calculates the VRRP checksum according to draft-ietf-vrrp-ipv6-spec-08.txt. Therefore, when **tcpdump** parses IPv6 VRRP packets that are received from older Junos OS releases (prior to Junos OS Release 12.2), the **bad vrrp cksum** message is displayed:

```
23:20:32.657328 Out
...
-----original packet-----
00:00:5e:00:02:03 > 33:33:00:00:00:12, ethertype IPv6 (0x86dd), length
94: (class 0xc0, hlim 255, next-header: VRRP (112), length: 40)
fe80::224:dcff:fe47:57f > ff02::12: VRRPv3-advertisement 40: vrid=3 prio=100
intvl=100(centisec) (bad vrrp cksum b4e2!) addrs(2):
fe80::200:5eff:fe00:3,2001:4818:f000:14::1
3333 0000 0012 0000 5e00 0203 86dd 6c00
0000 0028 70ff fe80 0000 0000 0000 0224
dcff fe47 057f ff02 0000 0000 0000 0000
0000 0000 0012 3103 6402 0064 b4e2 fe80
0000 0000 0000 0200 5eff fe00 0003 2001
4818 f000 0014 0000 0000 0000 0001
```

You can ignore this message because it does not indicate VRRP failure.

- When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements (for IPv4 VRRP) cannot be configured for any VRRP groups. Therefore, if authentication is required, you need to configure alternative non-VRRP authentication mechanisms.

- When VRRPv3 is enabled, the **advertise-interval** statement (for IPv4 VRRP) and the **inet6-advertise-interval** statement (for IPv6 VRRP) cannot be used to configure advertisement intervals. Instead, use the **fast-interval** statement to configure advertisement intervals.
- VRRPv3 for IPv4 does not interoperate with the previous versions of VRRP. If VRRPv2 IPv4 advertisement packets are received by a router on which VRRPv3 is enabled, the router transitions itself to the backup state to avoid creating multiple masters in the network. Due to this behavior, you must be cautious when enabling VRRPv3 on your existing VRRPv2 networks. See [“Understanding VRRPv2 to VRRPv3 Transition” on page 58](#) for more information.



NOTE: VRRPv3 advertisement packets are ignored by the routers on which previous versions of VRRP are configured.

Understanding VRRPv2 to VRRPv3 Transition

You must enable VRRPv3 in your network only if VRRPv3 can be enabled on all the VRRP routers in your network. Even if VRRPv3 can be enabled on all the VRRP routers in your network, care must be taken to avoid traffic loss when you transition your network to VRRPv3. This is because it is practically not possible to configure VRRPv3 on all routers simultaneously. There is a small time frame in the transition period during which VRRPv2 and VRRPv3 coexist in the network. During this period, to avoid having multiple masters in the network, the VRRPv3 IPv4 routers switch to the backup state when they receive a VRRPv2 IPv4 advertisement packet. VRRPv2 IPv4 packets are always given the highest priority. Additionally, to avoid having multiple masters in your IPv6 network due to checksum differences, you need to disable VRRP for IPv6 on the backup routers.



NOTE: Configuration change from VRRPv2 to VRRPv3 (or VRRPv3 to VRRPv2) restarts the VRRP state machine on all the configured VRRP groups.

The following example illustrates the steps and events that take place during a VRRPv2 to VRRPv3 transition:

Consider a scenario where two VRRPv2 routers, R1 and R2, are configured in two groups, G1 and G2. The R1 router acts as the master for G1 and the R2 router acts as the master for G2. [Table 17 on page 58](#) lists the transition steps and events for this setup:

Table 17: Example: VRRPv2 to VRRPv3 Transition Steps and Events

1. Upgrade the R1 router with Junos OS Release 12.2 or later.
 - R2 becomes master for both G1 and G2.
 - After the upgrade of the R1 router is completed, R1 becomes the master for G1. R2 remains as the master for G2.
 2. Upgrade the R2 router with Junos OS Release 12.2 or later.
 - R1 becomes master for both G1 and G2.
 - After the upgrade of R2 router is completed, R2 becomes the master for G2. R1 remains as the master for G1.
-

Table 17: Example: VRRPv2 to VRRPv3 Transition Steps and Events (*continued*)

For IPv4	For IPv6
3. Enable VRRPv3 on the R1 router. <ul style="list-style-type: none"> Because VRRPv2 IPv4 advertisement packets are given higher priority, R1 becomes the backup for both G1 and G2. 	3. Deactivate the G1 and G2 groups on the R2 router. <ul style="list-style-type: none"> G1 and G2 groups on the R1 router become master.
4. Enable VRRPv3 on the R2 router. <ul style="list-style-type: none"> R1 becomes the master for G1 and R2 becomes the master for G2. 	4. Enable VRRPv3 on the R1 router. <ul style="list-style-type: none"> R1 becomes master for both G1 and G2.
	5. Enable VRRPv3 on the R2 router.
	6. Activate G1 and G2 groups on the R2 router. <ul style="list-style-type: none"> R2 becomes master for G2. R1 remains as the master for G1.

Related Documentation

- Understanding High Availability Features on Juniper Networks Routers
- [Configuring Basic VRRP Support on page 180](#)
- VRRP Configuration Hierarchy
- VRRP for IPv6 Configuration Hierarchy
- [authentication-type on page 208](#)
- [authentication-key on page 207](#)
- [fast-interval on page 211](#)
- [inet6-advertise-interval on page 214](#)
- [version-3 on page 225](#)
- [virtual-link-local-address on page 227](#)

Improving the Convergence Time for VRRP

You can enable faster convergence time for the configured Virtual Router Redundancy Protocol (VRRP), thereby reducing the traffic restoration time to less than 1 second. To improve the convergence time for the VRRP, perform the following tasks:

- **Configure the distributed periodic packet management process**—When the VRRP process is busy and does not send VRRP advertisements, the backup VRRP routers might assume that the master router is down and take over as the master router, causing unnecessary flaps. To address this problem and to reduce the load on the VRRP process, Junos OS uses the distributed periodic packet management (PPM) process to send VRRP advertisements on behalf of the VRRP process.

To configure the distributed PPM process, include the **delegate-processing** statement at the **[edit protocols vrrp]** hierarchy level.

- **Disable the skew timer**—The skew timer in VRRP is used to ensure that two backup routers do not switch to the master state at the same time in case of a failover situation. When there is only one master router and one backup router in the network deployment,

you can disable the skew timer, thereby reducing the time required to transition to the master state.

To disable the skew timer, include the **skew-timer-disable** statement at the **[edit protocols vrrp]** hierarchy level.

- **Configure the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state**—The backup router waits until a certain number of advertisement packets are lost after which it transitions to the master state. This waiting time can be fatal in scenarios such as router failure or link failure. To avoid such a situation and to enable faster convergence time, in Junos OS Release 12.2 and later, you can configure a fast advertisement interval value that specifies the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state.

To configure the fast advertisement interval, include the **global-advertisements-threshold** statement at the **[edit protocols vrrp]** hierarchy level.

- **Configure inheritance of VRRP groups**—Junos OS enables you to configure VRRP groups on the various subnets of a virtual LAN (VLAN) to inherit the state and configuration of one of the groups, which is known as the active VRRP group. When the **vrrp-inherit-from** statement is included in the configuration, only the active VRRP group, from which the other VRRP groups inherit the state, sends out frequent VRRP advertisements and processes incoming VRRP advertisements. Use inherit groups for scaled configurations. For example, if you have 1000 VRRP groups with an advertisement interval of 100 ms, then use inherit groups.

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]** hierarchy level.



NOTE:

- The reduction in convergence time is not applicable when VRRP is configured over integrated routing and bridging (IRB) interfaces, aggregated Ethernet interfaces, and multichassis link aggregation group (MC-LAG) interfaces.
- Compared to other routers, the convergence time and the traffic restoration time are less for MX Series routers with MPCs.
- Reduction in convergence time is applicable for all types of configurations at the physical interface but the convergence time might not be less than 1 second for all the configurations. The convergence time depends on the number of groups that are transitioning from the backup to the master state and the interval at which these groups are transitioning.

Related
Documentation

- [Configuring Inheritance for a VRRP Group on page 192](#)
- [Configuring VRRP to Improve Convergence Time on page 196](#)
- [delegate-processing on page 210](#)

- [global-advertisements-threshold](#) on page 212
- [skew-timer-disable](#) on page 221

PART 2

Configuration

- [Configuration: GRES on page 65](#)
- [Configuration Statements: GRES on page 69](#)
- [Configuration: Graceful Restart on page 79](#)
- [Configuration Statements: Graceful Restart on page 119](#)
- [Configuration: NSB on page 133](#)
- [Configuration Statements: NSB on page 135](#)
- [Configuration: NSR on page 137](#)
- [Configuration Statements: NSR on page 147](#)
- [Configuration: Unified ISSU on page 153](#)
- [Configuration Statements: Unified ISSU on page 173](#)
- [Configuration: VRRP on page 179](#)
- [Configuration Statements: VRRP on page 203](#)

CHAPTER 7

Configuration: GRES

- [Configuring Graceful Routing Engine Switchover on page 65](#)
- [Resetting Local Statistics on page 66](#)

Configuring Graceful Routing Engine Switchover

This section contains the following topics:

- [Enabling Graceful Routing Engine Switchover on page 65](#)
- [Synchronizing the Routing Engine Configuration on page 65](#)
- [Verifying Graceful Routing Engine Switchover Operation on page 66](#)

Enabling Graceful Routing Engine Switchover

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

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

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

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

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

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 graceful Routing Engine switchover, 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 graceful Routing Engine switchover is enabled, on the backup Routing Engine, issue the **show system switchover** command. When the output of the command indicates that the **Graceful switchover** field is set to **on**, graceful Routing Engine switchover 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 Junos OS Operational Mode Commands.

Related Documentation

- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 3](#)
- [Graceful Routing Engine Switchover System Requirements on page 7](#)
- [Requirements for Routers with a Backup Router Configuration on page 10](#)
- [Resetting Local Statistics on page 66](#)
- [graceful-switchover on page 77](#)

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 Junos OS Operational Mode Commands.



.....

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 in the Junos OS on page 3](#)
- [Configuring Graceful Routing Engine Switchover on page 65](#)

CHAPTER 8

Configuration Statements: GRES

- [\[edit chassis\] Hierarchy Level on page 69](#)

[\[edit chassis\] Hierarchy Level](#)

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection {
          non-revertive;
        }
        system-priority;
      }
    }
    sonet {
      device-count number;
    }
    maximum-links maximum-links-limit;
  }
  alarm {
    ds1 {
      ais (ignore | red | yellow);
      ylw (ignore | red | yellow);
    }
    ethernet {
      link-down (ignore | red | yellow);
    }
    integrated-services {
      failure (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
    relay
    input {
      port port-number {
        mode (close | open);
        trigger (ignore | red | yellow);
      }
    }
  }
}
```

```

output{
    port port-number {
        input-relay input-relay;
        mode (close | open);
        temperature;
    }
}
serial {
    cts-absent (ignore | red | yellow);
    dcd-absent (ignore | red | yellow);
    dsr-absent (ignore | red | yellow);
    loss-of-rx-clock (ignore | red | yellow);
    loss-of-tx-clock (ignore | red | yellow);
    tm-absent (ignore | red | yellow);
}
services {
    hw-down (ignore | red | yellow);
    linkdown (ignore | red | yellow);
    pic-hold-reset (ignore | red | yellow);
    pic-reset (ignore | red | yellow);
    rx-errors (ignore | red | yellow);
    sw-down (ignore | red | yellow);
    tx-errors (ignore | red | yellow);
}
sonet {
    (ais-l | ais-p | ber-sd | ber-sf | locd | lol | lop-p | los | pll | plm-p | rfi-l | rfl-p | uneq-p)
    (ignore | red | yellow);
}
t3 {
    (ais | exz | ferf | idle | lcv | lof | los | pll | ylw) (ignore | red | yellow);
}
}
cluster {
    control-link-recovery;
    control-ports {
        fpc slot-number port port-number;
    }
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    redundancy-group {
        ... the redundancy-group subhierarchy appears at the end of the [edit chassis cluster]
        hierarchy ...
    }
}
reth-count number;
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag;
    level severity;
    no-remote-trace;
}
redundancy-group group-number {
    gratuitous-arp-count number;
    hold-down-interval seconds;
    interface-monitor {
        interface-name weight number;
    }
}

```



```

    }
    ip-monitoring {
        family {
            inet {
                ipv4-address {
                    interface rethindex.logical-unit-number secondary-ip-address ipv4-address;
                    weight number;
                }
            }
        }
        global-threshold number;
        global-weight number;
        retry-count count;
        retry-interval interval;
    }
    node node-number priority priority-number;
    preempt;
}
config-button {
    no-clear;
    no-rescue;
}
container-devices {
    device-count number;
}
craft-lockout;
disable-power-management;
disk-partition partition-name (/config | /var) {
    level (full | high) {
        free-space threshold-value (mb | percent);
    }
}
enhanced-policer;
extended-statistics;
fabric {
    degraded {
        action-fpc-restart-disable;
        degraded-fabric-detection-enable
        degraded-fpc-bad-plane-threshold number-bad-planes;
    }
    redundancy-mode (increased-bandwidth | redundant);
}
filter;
fpc slot-number {
    ... the fpc subhierarchy appears after the main [edit chassis] hierarchy ...
}
fpc-feb-connectivity {
    fpc slot-number feb (slot-number | none);
}
fpc-resync;
fru-poweron-sequence sequence;
lcc index {
    ... the lcc subhierarchy appears after the main [edit chassis] hierarchy ...
}
maximum-ecmp value;
memory-enhanced {

```

```
filter;
route;
vpn-label;
}
network-services (ethernet | enhanced-ethernet | ip | enhanced-ip);
(packet-scheduling | no-packet-scheduling);
pem {
    minimum number;
}
policer-drop-probability-low;
ppp-subscriber-services (disable | enable);
redundancy {
    cfeb slot (always | preferred);
    failover {
        on-disk-failure;
        on-loss-of-keepalives;
    }
    feb {
        redundancy-group group-name {
            description description;
            feb slot-number <backup | primary>;
            no-auto-failover;
        }
    }
    graceful-switchover;
    keepalive-time seconds;
    routing-engine slot-number (backup | disabled | master);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
}
route-memory-enhanced;
route-localization {
    inet (chassis);
    inet6;
}
routing-engine {
    bios {
        no-auto-upgrade;
    }
    on-disk-failure disk-failure-action (halt | reboot);
}
sfm slot-number {
    power off;
}
sib {
    minimum number;
}
(source-route | no-source-route);
state [
    cb-upgrade [on | off];
]
synchronization { # for M Series and T Series routers
    primary (external-a | external-b);
    secondary (external-a | external-b);
    signal-type (e1 | t1);
    switching-mode (non-revertive | revertive);
```

```

transmitter-enable;
validation-interval seconds;
y-cable-line-termination;
}
synchronization { # for MX80 and MX240 routers
  clock-mode (auto-select | free-run);
  esmc-transmit {
    interfaces (all | interface-name);
  }
  hold-interval {
    configuration-change seconds;
    restart seconds;
    switchover seconds;
  }
  network-option (option-1 | option-2);
  quality-mode-enable;
  selection-mode (configured-quality|received-quality);
  source {
    (external-a | external-b) {
      priority number;
      quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
      request (force-switch | lockout);
    }
    interfaces interface-name {
      priority number;
      quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
      request (force-switch | lockout);
      wait-to-restore minutes;
    }
  }
  switchover-mode (revertive | non-revertive);
}
synchronization { # for ACX Series routers
  clock-mode (auto-select | free-run);
  esmc-transmit {
    interfaces (all | interface-name);
  }
  hold-interval {
    configuration-change seconds;
    restart seconds;
    switchover seconds;
  }
  network-option (option-1 | option-2);
  quality-mode-enable;
  selection-mode (configured-quality | received-quality);
  source {
    (bits | gps) {
      priority number;
      quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
      request (force-switch | lockout);
    }
    interfaces interface-name {
      priority number;
      quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
      request (force-switch | lockout);
      wait-to-restore minutes;
    }
  }
}

```

```

    }
  }
  switchover-mode(non-revertive | revertive);
}
system-domains {
  protected-system-domains psdnumerical-index {
    control-plane-bandwidth-percent percent;
    control-slot-numbers [ slot-numbers ];
    control-system-id control-system-id;
    description description;
    fpcs [ slot-numbers ];
  }
  root-domain-id root-domain-id;
}
vrf-mtu-check;
}

chassis {
  fpc slot-number {
    number-of-ports active-ports;
    offline;
    pic slot-number {
      ... the pic subhierarchy appears after the main [edit chassis fpc slot-number] hierarchy
      ...
    }
    port-mirror-instance port-mirror-instance-name;
    power (off | on);
    sampling-instance instance-name;
  }

  fpc slot-number {
    pic slot-number {
      adaptive-services {
        service-package (layer-2 | layer-3 | ...the following extension-provider subhierarchy
          ...);
        extension-provider {
          control-cores number;
          data-cores number;
          data-flow-affinity {
            hash-key (layer-3 | layer-4);
          }
          channelization;
          forwarding-db-size megabytes;
          object-cache-size megabytes;
          package package-name;
          policy-db-size megabytes;
          syslog {
            facility {
              severity;
              destination (pic-console | routing-engine);
            }
          }
          wired-process-mem-size megabytes;
        }
      }
    }
  }
  aggregated-devices {

```

```

    ima {
        device-count number;
    }
}
aggregate-ports;
atm-cell-relay-accumulation;
atm-l2circuit-mode (aal5 | cell | trunk trunk);
cel {
    e1 port-number {
        channel-group group-number timeslots slot-number;
    }
}
ct3 {
    port port-number {
        t1 link-number {
            channel-group group-number timeslots slot-number;
        }
    }
}
ethernet {
    pic-mode (enhanced-switching | routing | switching);
}
fibre-channel {
    port port-number;
    port-range port-range-low port-range-high
}
egress-policer-overhead bytes;
forwarding-mode {
    sa-multicast;
    vlan-steering {
        vlan-rule (high-low | odd-even);
    }
}
framing (e1 | e3 | sdh | sonet | t1 | t3);
idle-cell-format {
    itu-t;
    payload-pattern payload-pattern-byte;
}
ingress-policer-overhead bytes;
inline-services {
    bandwidth (1g | 10g);
}
linerate-mode;
max-queues-per-interface (4 | 8);
mlfr-uni-nni-bundles number;
no-concatenate;
no-multi-rate;
port port-number {
    framing (e1 | e3 | sdh | sonet | t1 | t3);
    forwarding-mode {
        sa-multicast;
    }
    speed ( oc3-stm1 | oc12-stm4 | oc48-stm16);
}
port-mirror-instance port-mirror-instance-name;
q-pic-large-buffer {

```

```

        (large-scale | small-scale);
    }
    red-buffer-occupancy {
        weighted-averaged <instant-usage-weight-exponent weight-value>;
    }
    shdsl {
        pic-mode (1-port-atm | 2-port-atm);
    }
    sparse-dlcis;
    traffic-manager {
        egress-shaping-overhead number;
        ingress-shaping-overhead number;
        mode {
            egress-only;
            ingress-and-egress;
            session-shaping;
        }
    }
    tunnel-queuing;
    tunnel-services {
        bandwidth (1g | 10g | 20g | 40g);
        tunnel-only;
    }
    vtmapping (itu-t | klm);
}

chassis {
    lcc index {
        fpc slot-number {
            ... the fpc subhierarchy appears after the main [edit chassis lcc index] hierarchy ...
        }
        offline;
        online-expected;
    }
}

lcc index {
    fpc slot-number {
        pic slot-number {
            ... the pic subhierarchy appears after the main [edit chassis lcc index fpc slot-number] hierarchy ...
        }
        power (off | on);
        sampling-instance instance-name;
    }

    fpc slot-number {
        pic slot-number {
            aggregate-ports;
            atm-cell-relay-accumulation;
            atm-l2circuit-mode (aal5 | cell | trunk trunk);
            framing (e1 | e3 | sdh | sonet | t1 | t3);
            idle-cell-format {
                itu-t;
            }
        }
    }
}

```

```

        payload-pattern payload-pattern-byte;
    }
    linerate-mode;
    max-queues-per-interface (4 | 8);
    no-concatenate;
    no-mcast-replication;
    no-pre-classifier;
    port port-number {
        framing (e1 | e3 | sdh | sonet | t1 | t3);
    }
    q-pic-large-buffer {
        (large-scale | small-scale);
    }
    red-buffer-occupancy {
        weighted-averaged <instant-usage-weight-exponent weight-value>;
    }
    shdsl {
        pic-mode (1-port-atm | 2-port-atm);
    }
    traffic-manager {
        egress-shaping-overhead bytes;
        ingress-shaping-overhead bytes;
        mode {
            egress-only;
            ingress-and-egress;
        }
    }
}
}
}

```

Related Documentation • [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

graceful-switchover

Syntax	<code>graceful-switchover;</code>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Configuring Graceful Routing Engine Switchover on page 65

Configuration: Graceful Restart

- [Enabling Graceful Restart on page 79](#)
- [Configuring Routing Protocols Graceful Restart on page 80](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 86](#)
- [Configuring VPN Graceful Restart on page 88](#)
- [Configuring Logical System Graceful Restart on page 89](#)
- [Example: Configuring Graceful Restart on page 91](#)
- [Example: Managing Helper Modes for OSPF Graceful Restart on page 116](#)

Enabling Graceful Restart

Graceful restart is disabled by default. You must configure graceful restart at the **[edit routing-options]** hierarchy level to enable the feature.

For graceful restart to function properly, graceful restart must be enabled on the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

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



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.

To disable graceful restart, include the **disable** statement. To configure a time period for complete restart, include the **restart-duration** statement. You can specify a number between 120 and 900.

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

When you include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level, graceful restart is also enabled for aggregate and static routes.

- Related Documentation**
- [Graceful Restart Concepts on page 29](#)
 - [Graceful Restart System Requirements on page 30](#)
 - [Aggregate and Static Routes on page 31](#)
 - [Example: Configuring Graceful Restart on page 91](#)

Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- [Enabling Graceful Restart on page 80](#)
- [Configuring Graceful Restart Options for BGP on page 81](#)
- [Configuring Graceful Restart Options for ES-IS on page 82](#)
- [Configuring Graceful Restart Options for IS-IS on page 82](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 83](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 84](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 84](#)
- [Tracking Graceful Restart Events on page 86](#)

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: 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 86.

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 | ospf3) 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 | ospf3) 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 | ospf3) graceful-restart]` hierarchy level.

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

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

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

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

To 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 86](#).



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 29](#)
- [Graceful Restart System Requirements on page 30](#)
- [Graceful Restart and Routing Protocols on page 31](#)
- [Verifying Graceful Restart Operation on page 234](#)
- [Example: Configuring Graceful Restart on page 91](#)

Configuring Graceful Restart for MPLS-Related Protocols

This section contains the following topics:

- [Configuring Graceful Restart Globally on page 86](#)
- [Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 87](#)
- [Configuring Graceful Restart Options for LDP on page 87](#)

Configuring Graceful Restart Globally

To configure graceful restart globally for all MPLS-related protocols, include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level:

```
[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.

Configuring Graceful Restart Options for RSVP, CCC, and TCC

Because CCC and TCC rely on RSVP, you must modify these three protocols as a single group.

To configure how long the router retains the state of its RSVP neighbors while they undergo a graceful restart, include the **maximum-helper-recovery-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to recover.

To configure the delay between when the router discovers that a neighboring router has gone down and when it declares the neighbor down, include the **maximum-helper-restart-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to restart.

```
[edit]
protocols {
  rsvp {
    graceful-restart {
      disable;
      helper-disable;
      maximum-helper-recovery-time;
      maximum-helper-restart-time;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable RSVP, CCC, and TCC graceful restart, include the **disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. To disable RSVP, CCC, and TCC helper capability, include the **helper-disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for LDP

When configuring graceful restart for LDP, you can include the following optional statements at the **[edit protocols ldp graceful-restart]** hierarchy level:

```
[edit protocols ldp graceful-restart]
disable;
helper-disable;
maximum-neighbor-reconnect-time seconds;
maximum-neighbor-recovery-time seconds;
reconnect-time seconds;
recovery-time seconds;

[edit routing-options]
graceful-restart;
```

The statements have the following effects on the graceful restart process:

- To configure the length of time required to reestablish a session after a graceful restart, include the **reconnect-time** statement; the range is 30 through 300 seconds. To limit the maximum reconnect time allowed from a restarting neighbor router, include the **maximum-neighbor-reconnect-time** statement; the range is 30 through 300 seconds.
- To configure the length of time that helper routers are required to maintain the old forwarding state during a graceful restart, include the **recovery-time** statement; the range is 120 through 1800 seconds. On the helper router, you can configure a statement that overrides the request from the restarting router and sets the maximum length of time the helper router will maintain the old forwarding state. To configure this feature, include the **maximum-neighbor-recovery-time** statement; the range is 140 through 1900 seconds.



NOTE: The value for the **recovery-time** and **maximum-neighbor-recovery-time** statements at the **[edit protocols ldp graceful-restart]** hierarchy level should be approximately 80 seconds longer than the value for the **restart-duration** statement at the **[edit routing-options graceful-restart]** hierarchy level. Otherwise, a warning message appears when you try to commit the configuration.

- To disable LDP graceful restart capability, include the **disable** statement. To disable LDP graceful restart helper capability, include the **helper-disable** statement.

Configuring VPN Graceful Restart

Graceful restart allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. Without graceful restart, a control plane restart disrupts any VPN services provided by the router. Graceful restart is supported on Layer 2 VPNs, Layer 3 VPNs, virtual-router routing instances, and VPLS.

To implement graceful restart for a Layer 2 VPN or Layer 3 VPN, perform the configuration tasks described in the following sections:

- [Configuring Graceful Restart Globally on page 88](#)
- [Configuring Graceful Restart for the Routing Instance on page 89](#)

Configuring Graceful Restart Globally

To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To configure a global duration for the graceful restart period, include the **restart-duration** statement at the **[edit routing-options graceful-restart]** hierarchy level.

```
[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.

Configuring Graceful Restart for the Routing Instance

For Layer 3 VPNs only, you must also configure graceful restart for all routing and MPLS-related protocols within a routing instance by including the **graceful-restart** statement at the **[edit routing-instances *instance-name* routing-options]** hierarchy level. Because you can configure multi-instance BGP and multi-instance LDP, graceful restart for a carrier-of-carriers scenario is supported. To configure the duration of the graceful restart period for the routing instance, include the **restart-duration** statement at the **[edit routing-instances *instance-name* routing-options]**.

```

[edit]
routing-instances {
  instance-name {
    routing-options {
      graceful-restart {
        disable;
        restart-duration seconds;
      }
    }
  }
}

```

You can disable graceful restart for individual protocols with the **disable** statement at the **[edit routing-instances *instance-name* protocols *protocol-name* graceful-restart]** hierarchy level.

Related Documentation

- [Graceful Restart Concepts on page 29](#)
- [Graceful Restart System Requirements on page 30](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 35](#)
- [Verifying Graceful Restart Operation on page 234](#)
- [Example: Configuring Graceful Restart on page 91](#)

Configuring Logical System Graceful Restart

Graceful restart for a logical system functions much as graceful restart does in the main router. The only difference is the location of the **graceful-restart** statement.

The following topics describe what to configure to implement graceful restart in a logical system:

- [Enabling Graceful Restart Globally on page 90](#)
- [Configuring Graceful Restart for a Routing Instance on page 90](#)

Enabling Graceful Restart Globally

To enable graceful restart in a logical system, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-options]** hierarchy level. To configure a global duration of the graceful restart period, include the **restart-duration** statement at the **[edit logical-systems *logical-system-name* routing-options graceful-restart]** hierarchy level.

```
[edit]
logical-systems {
  logical-system-name {
    routing-options {
      graceful-restart {
        disable;
        restart-duration seconds;
      }
    }
  }
}
```

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

Configuring Graceful Restart for a Routing Instance

For Layer 3 VPNs only, you must also configure graceful restart globally for a routing instance inside a logical system. To configure, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]** hierarchy level. Because you can configure multi-instance BGP and multi-instance LDP, graceful restart for a carrier-of-carriers scenario is supported. To configure the duration of the graceful restart period for the routing instance, include the **restart-duration** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]**.

```
[edit]
logical-systems {
  logical-system-name {
    routing-instances {
      instance-name {
        routing-options {
          graceful-restart {
            disable;
            restart-duration seconds;
          }
        }
      }
    }
  }
}
```

To disable graceful restart for individual protocols with the **disable** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols *protocol-name* graceful-restart]** hierarchy level.

- Related Documentation**
- [Graceful Restart Concepts on page 29](#)
 - [Graceful Restart System Requirements on page 30](#)
 - [Graceful Restart on Logical Systems on page 36](#)
 - [Verifying Graceful Restart Operation on page 234](#)
 - [Example: Configuring Graceful Restart on page 91](#)

Example: Configuring Graceful Restart

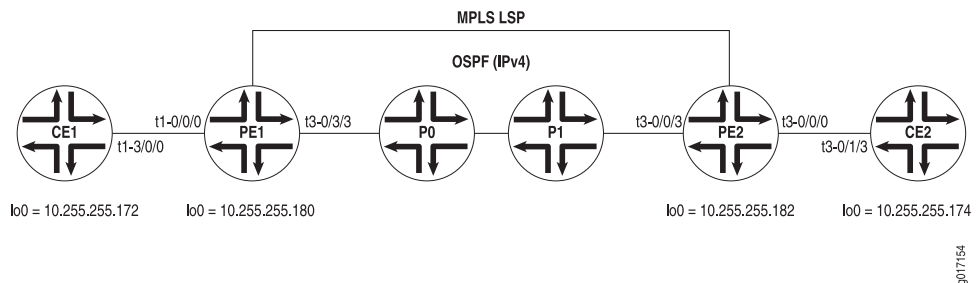
To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level as well as in the protocol level.

For example:

```
protocols {
  bgp {
    group ext {
      graceful-restart;
    }
  }
}
routing-options {
  graceful-restart;
}
```

Figure 8 on page 91 shows a standard MPLS VPN network. Routers CE1 and CE2 are customer edge routers, PE1 and PE2 are provider edge routers, and P0 is a provider core router. Several Layer 3 VPNs are configured across this network, as well as one Layer 2 VPN. Interfaces are shown in the diagram and are not included in the configuration example that follows.

Figure 8: Layer 3 VPN Graceful Restart Topology



Router CE1 On Router CE1, configure the following protocols on the logical interfaces of **t3-3/1/0**: OSPF on unit 101, RIP on unit 102, BGP on unit 103, and IS-IS on unit 512. Also configure graceful restart, BGP, IS-IS, OSPF, and RIP on the main instance to be able to connect to the routing instances on Router PE1.

```
[edit]
interfaces {
  t3-3/1/0 {
```

```
encapsulation frame-relay;
unit 100 {
    dlci 100;
    family inet {
        address 10.96.100.2/30;
    }
}
unit 101 {
    dlci 101;
    family inet {
        address 10.96.101.2/30;
    }
}
unit 102 {
    dlci 102;
    family inet {
        address 10.96.102.2/30;
    }
}
unit 103 {
    dlci 103;
    family inet {
        address 10.96.103.2/30;
    }
}
unit 512 {
    dlci 512;
    family inet {
        address 10.96.252.1/30;
    }
}
}
lo0 {
    unit 0 {
        family inet {
            address 10.245.14.172/32;
            primary;
        }
        address 10.96.110.1/32;
        address 10.96.111.1/32;
        address 10.96.112.1/32;
        address 10.96.113.1/32;
        address 10.96.116.1/32;
    }
    family iso {
        address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4172.00;
    }
}
}
routing-options {
    graceful-restart;
    autonomous-system 65100;
}
protocols {
    bgp {
        group CE-PE-INET {
```

```

        type external;
        export BGP_INET_LB_DIRECT;
        neighbor 10.96.103.1 {
            local-address 10.96.103.2;
            family inet {
                unicast;
            }
            peer-as 65103;
        }
    }
}
isis {
    export ISIS_L2VPN_LB_DIRECT;
    interface t3-3/1/0.512;
}
ospf {
    export OSPF_LB_DIRECT;
    area 0.0.0.0 {
        interface t3-3/1/0.101;
    }
}
rip {
    group RIP {
        export RIP_LB_DIRECT;
        neighbor t3-3/1/0.102;
    }
}
}
policy-options {
    policy-statement OSPF_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.101.0/30 exact;
                route-filter 10.96.111.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
    policy-statement RIP_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.102.0/30 exact;
                route-filter 10.96.112.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
    policy-statement BGP_INET_LB_DIRECT {

```

```
term direct {
  from {
    protocol direct;
    route-filter 10.96.103.0/30 exact;
    route-filter 10.96.113.1/32 exact;
  }
  then accept;
}
term final {
  then reject;
}
}
policy-statement ISIS_L2VPN_LB_DIRECT {
  term direct {
    from {
      protocol direct;
      route-filter 10.96.116.1/32 exact;
    }
    then accept;
  }
  term final {
    then reject;
  }
}
}
```

Router PE1 On Router PE1, configure graceful restart in the master instance, along with BGP, OSPF, MPLS, and LDP. Next, configure several protocol-specific instances of graceful restart. By including instances for BGP, OSPF, Layer 2 VPNs, RIP, and static routes, you can observe the wide range of options available when you implement graceful restart. Configure the following protocols in individual instances on the logical interfaces of **t3-0/0/0**: a static route on unit 100, OSPF on unit 101, RIP on unit 102, BGP on unit 103, and Frame Relay on unit 512 for the Layer 2 VPN instance.

```
[edit]
interfaces {
  t3-0/0/0 {
    dce;
    encapsulation frame-relay-ccc;
    unit 100 {
      dlci 100;
      family inet {
        address 10.96.100.1/30;
      }
      family mpls;
    }
    unit 101 {
      dlci 101;
      family inet {
        address 10.96.101.1/30;
      }
      family mpls;
    }
    unit 102 {
      dlci 102;
      family inet {
```



```

        address 10.96.102.1/30;
    }
    family mpls;
}
unit 103 {
    dlci 103;
    family inet {
        address 10.96.103.1/30;
    }
    family mpls;
}
unit 512 {
    encapsulation frame-relay-ccc;
    dlci 512;
}
}
t1-0/1/0 {
    unit 0 {
        family inet {
            address 10.96.0.2/30;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.245.14.176/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4176.00;
        }
    }
}
}
routing-options {
    graceful-restart;
    router-id 10.245.14.176;
    autonomous-system 69;
}
protocols {
    mpls {
        interface all;
    }
    bgp {
        group PEPE {
            type internal;
            neighbor 10.245.14.182 {
                local-address 10.245.14.176;
                family inet-vpn {
                    unicast;
                }
                family l2vpn {
                    unicast;
                }
            }
        }
    }
}

```

```
    }
  }
  ospf {
    area 0.0.0.0 {
      interface t1-0/1/0.0;
      interface fxp0.0 {
        disable;
      }
      interface lo0.0 {
        passive;
      }
    }
  }
  ldp {
    interface all;
  }
}
policy-options {
  policy-statement STATIC-import {
    from community STATIC;
    then accept;
  }
  policy-statement STATIC-export {
    then {
      community add STATIC;
      accept;
    }
  }
  policy-statement OSPF-import {
    from community OSPF;
    then accept;
  }
  policy-statement OSPF-export {
    then {
      community add OSPF;
      accept;
    }
  }
  policy-statement RIP-import {
    from community RIP;
    then accept;
  }
  policy-statement RIP-export {
    then {
      community add RIP;
      accept;
    }
  }
  policy-statement BGP-INET-import {
    from community BGP-INET;
    then accept;
  }
  policy-statement BGP-INET-export {
    then {
      community add BGP-INET;
      accept;
    }
  }
}
```

```

    }
  }
  policy-statement L2VPN-import {
    from community L2VPN;
    then accept;
  }
  policy-statement L2VPN-export {
    then {
      community add L2VPN;
      accept;
    }
  }
  community BGP-INET members target:69:103;
  community L2VPN members target:69:512;
  community OSPF members target:69:101;
  community RIP members target:69:102;
  community STATIC members target:69:100;
}
routing-instances {
  BGP-INET {
    instance-type vrf;
    interface t3-0/0/0.103;
    route-distinguisher 10.245.14.176:103;
    vrf-import BGP-INET-import;
    vrf-export BGP-INET-export;
    routing-options {
      graceful-restart;
      autonomous-system 65103;
    }
    protocols {
      bgp {
        group BGP-INET {
          type external;
          export BGP-INET-import;
          neighbor 10.96.103.2 {
            local-address 10.96.103.1;
            family inet {
              unicast;
            }
          }
          peer-as 65100;
        }
      }
    }
  }
}
L2VPN {
  instance-type l2vpn;
  interface t3-0/0/0.512;
  route-distinguisher 10.245.14.176:512;
  vrf-import L2VPN-import;
  vrf-export L2VPN-export;
  protocols {# There is no graceful-restart statement for Layer 2 VPN instances.
    l2vpn {
      encapsulation-type frame-relay;
      site CE1-ISIS {
        site-identifier 512;
      }
    }
  }
}

```

```
        interface t3-0/0/0.512 {
            remote-site-id 612;
        }
    }
}
}
OSPF {
    instance-type vrf;
    interface t3-0/0/0.101;
    route-distinguisher 10.245.14.176:101;
    vrf-import OSPF-import;
    vrf-export OSPF-export;
    routing-options {
        graceful-restart;
    }
    protocols {
        ospf {
            export OSPF-import;
            area 0.0.0.0 {
                interface all;
            }
        }
    }
}
RIP {
    instance-type vrf;
    interface t3-0/0/0.102;
    route-distinguisher 10.245.14.176:102;
    vrf-import RIP-import;
    vrf-export RIP-export;
    routing-options {
        graceful-restart;
    }
    protocols {
        rip {
            group RIP {
                export RIP-import;
                neighbor t3-0/0/0.102;
            }
        }
    }
}
STATIC {
    instance-type vrf;
    interface t3-0/0/0.100;
    route-distinguisher 10.245.14.176:100;
    vrf-import STATIC-import;
    vrf-export STATIC-export;
    routing-options {
        graceful-restart;
        static {
            route 10.96.110.1/32 next-hop t3-0/0/0.100;
        }
    }
}
```

```
}
```

Router P0 On Router P0, configure graceful restart in the main instance, along with OSPF, MPLS, and LDP. This allows the protocols on the PE routers to reach one another.

```
[edit]
interfaces {
  t3-0/1/3 {
    unit 0 {
      family inet {
        address 10.96.0.5/30;
      }
      family mpls;
    }
  }
  t1-0/2/0 {
    unit 0 {
      family inet {
        address 10.96.0.1/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.245.14.174/32;
      }
      family iso {
        address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4174.00;
      }
    }
  }
}
routing-options {
  graceful-restart;
  router-id 10.245.14.174;
  autonomous-system 69;
}
protocols {
  mpls {
    interface all;
  }
  ospf {
    area 0.0.0.0 {
      interface t1-0/2/0.0;
      interface t3-0/1/3.0;
      interface fxp0.0 {
        disable;
      }
      interface lo0.0 {
        passive;
      }
    }
  }
  ldp {
```

```
        interface all;  
    }  
}
```

Router PE2 On Router PE2, configure BGP, OSPF, MPLS, LDP, and graceful restart in the master instance. Configure the following protocols in individual instances on the logical interfaces of **t1-0/1/3**: a static route on unit 200, OSPF on unit 201, RIP on unit 202, BGP on unit 203, and Frame Relay on unit 612 for the Layer 2 VPN instance. Also configure protocol-specific graceful restart in all routing instances, except the Layer 2 VPN instance.

```
[edit]  
interfaces {  
    t3-0/0/0 {  
        unit 0 {  
            family inet {  
                address 10.96.0.6/30;  
            }  
            family mpls;  
        }  
    }  
    t1-0/1/3 {  
        dce;  
        encapsulation frame-relay-ccc;  
        unit 200 {  
            dlci 200;  
            family inet {  
                address 10.96.200.1/30;  
            }  
            family mpls;  
        }  
        unit 201 {  
            dlci 201;  
            family inet {  
                address 10.96.201.1/30;  
            }  
            family mpls;  
        }  
        unit 202 {  
            dlci 202;  
            family inet {  
                address 10.96.202.1/30;  
            }  
            family mpls;  
        }  
        unit 203 {  
            dlci 203;  
            family inet {  
                address 10.96.203.1/30;  
            }  
            family mpls;  
        }  
        unit 612 {  
            encapsulation frame-relay-ccc;  
            dlci 612;  
        }  
    }  
}
```

```

lo0 {
  unit 0 {
    family inet {
      address 10.245.14.182/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4182.00;
    }
  }
}
routing-options {
  graceful-restart;
  router-id 10.245.14.182;
  autonomous-system 69;
}
protocols {
  mpls {
    interface all;
  }
  bgp {
    group PEPE {
      type internal;
      neighbor 10.245.14.176 {
        local-address 10.245.14.182;
        family inet-vpn {
          unicast;
        }
        family l2vpn {
          unicast;
        }
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface t3-0/0/0.0;
    interface fxp0.0 {
      disable;
    }
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface all;
}
policy-options {
  policy-statement STATIC-import {
    from community STATIC;
    then accept;
  }
  policy-statement STATIC-export {
    then {
      community add STATIC;
    }
  }
}

```

```
        accept;
    }
}
policy-statement OSPF-import {
    from community OSPF;
    then accept;
}
policy-statement OSPF-export {
    then {
        community add OSPF;
        accept;
    }
}
policy-statement RIP-import {
    from community RIP;
    then accept;
}
policy-statement RIP-export {
    then {
        community add RIP;
        accept;
    }
}
policy-statement BGP-INET-import {
    from community BGP-INET;
    then accept;
}
policy-statement BGP-INET-export {
    then {
        community add BGP-INET;
        accept;
    }
}
policy-statement L2VPN-import {
    from community L2VPN;
    then accept;
}
policy-statement L2VPN-export {
    then {
        community add L2VPN;
        accept;
    }
}
community BGP-INET members target:69:103;
community L2VPN members target:69:512;
community OSPF members target:69:101;
community RIP members target:69:102;
community STATIC members target:69:100;
}
routing-instances {
    BGP-INET {
        instance-type vrf;
        interface t1-0/1/3.203;
        route-distinguisher 10.245.14.182:203;
        vrf-import BGP-INET-import;
        vrf-export BGP-INET-export;
    }
}
```



```

routing-options {
  graceful-restart;
  autonomous-system 65203;
}
protocols {
  bgp {
    group BGP-INET {
      type external;
      export BGP-INET-import;
      neighbor 10.96.203.2 {
        local-address 10.96.203.1;
        family inet {
          unicast;
        }
      }
      peer-as 65200;
    }
  }
}
}
L2VPN {
  instance-type l2vpn;
  interface t1-0/1/3.612;
  route-distinguisher 10.245.14.182:612;
  vrf-import L2VPN-import;
  vrf-export L2VPN-export;
  protocols {# There is no graceful-restart statement for Layer 2 VPN instances.
    l2vpn {
      encapsulation-type frame-relay;
      site CE2-ISIS {
        site-identifier 612;
        interface t1-0/1/3.612 {
          remote-site-id 512;
        }
      }
    }
  }
}
}
OSPF {
  instance-type vrf;
  interface t1-0/1/3.201;
  route-distinguisher 10.245.14.182:201;
  vrf-import OSPF-import;
  vrf-export OSPF-export;
  routing-options {
    graceful-restart;
  }
  protocols {
    ospf {
      export OSPF-import;
      area 0.0.0.0 {
        interface all;
      }
    }
  }
}
}

```

```
RIP {
  instance-type vrf;
  interface t1-0/1/3.202;
  route-distinguisher 10.245.14.182:202;
  vrf-import RIP-import;
  vrf-export RIP-export;
  routing-options {
    graceful-restart;
  }
  protocols {
    rip {
      group RIP {
        export RIP-import;
        neighbor t1-0/1/3.202;
      }
    }
  }
}
STATIC {
  instance-type vrf;
  interface t1-0/1/3.200;
  route-distinguisher 10.245.14.182:200;
  vrf-import STATIC-import;
  vrf-export STATIC-export;
  routing-options {
    graceful-restart;
  }
  static {
    route 10.96.210.1/32 next-hop t1-0/1/3.200;
  }
}
}
```

Router CE2 On Router CE2, complete the Layer 2 and Layer 3 VPN configuration by mirroring the protocols already set on Routers PE2 and CE1. Specifically, configure the following on the logical interfaces of **t1-0/0/3**: OSPF on unit 201, RIP on unit 202, BGP on unit 203, and IS-IS on unit 612. Finally, configure graceful restart, BGP, IS-IS, OSPF, and RIP on the main instance to be able to connect to the routing instances on Router PE2.

```
[edit]
interfaces {
  t1-0/0/3 {
    encapsulation frame-relay;
    unit 200 {
      dlci 200;
      family inet {
        address 10.96.200.2/30;
      }
    }
    unit 201 {
      dlci 201;
      family inet {
        address 10.96.201.2/30;
      }
    }
  }
}
```

```

unit 202 {
    dlci 202;
    family inet {
        address 10.96.202.2/30;
    }
}
unit 203 {
    dlci 203;
    family inet {
        address 10.96.203.2/30;
    }
}
unit 512 {
    dlci 512;
    family inet {
        address 10.96.252.2/30;
    }
}
}
lo0 {
    unit 0 {
        family inet {
            address 10.245.14.180/32 {
                primary;
            }
            address 10.96.210.1/32;
            address 10.96.111.1/32;
            address 10.96.212.1/32;
            address 10.96.213.1/32;
            address 10.96.216.1/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4180.00;
        }
    }
}
}
routing-options {
    graceful-restart;
    autonomous-system 65200;
}
protocols {
    bgp {
        group CE-PE-INET {
            type external;
            export BGP_INET_LB_DIRECT;
            neighbor 10.96.203.1 {
                local-address 10.96.203.2;
                family inet {
                    unicast;
                }
            }
            peer-as 65203;
        }
    }
}
}
isis {

```

```
export ISIS_L2VPN_LB_DIRECT;
interface t1-0/0/3.612;
}
ospf {
  export OSPF_LB_DIRECT;
  area 0.0.0.0 {
    interface t1-0/0/3.201;
  }
}
rip {
  group RIP {
    export RIP_LB_DIRECT;
    neighbor t1-0/0/3.202;
  }
}
}
policy-options {
  policy-statement OSPF_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.201.0/30 exact;
        route-filter 10.96.211.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement RIP_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.202.0/30 exact;
        route-filter 10.96.212.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement BGP_INET_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.203.0/30 exact;
        route-filter 10.96.213.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
}
```

```

policy-statement ISIS_L2VPN_LB_DIRECT {
  term direct {
    from {
      protocol direct;
      route-filter 10.96.216.1/32 exact;
    }
    then accept;
  }
  term final {
    then reject;
  }
}
}

```

Router PE1 Status Before a Restart The following example displays neighbor relationships on Router PE1 before a restart happens:

```

user@PE1> show bgp neighbor
Peer: 10.96.103.2+3785 AS 65100 Local: 10.96.103.1+179 AS 65103
  Type: External   State: Established   Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.96.110.1      Local ID: 10.96.103.1      Active Holdtime: 90
  Keepalive Interval: 30
  Local Interface: t3-0/0/0.103
  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: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI peer can save forwarding state: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table BGP-INET.inet.0 Bit: 30001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        0
    Suppressed due to damping: 0
  Last traffic (seconds): Received 8    Sent 3    Checked 3
  Input messages:  Total 15    Updates 0    Refreshes 0    Octets 321
  Output messages: Total 18    Updates 2    Refreshes 0    Octets 450
  Output Queue[2]: 0

Peer: 10.245.14.182+4701 AS 69   Local: 10.245.14.176+179 AS 69
  Type: Internal   State: Established   Flags: <>
  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: 10.245.14.176 Holdtime: 90 Preference: 170
Number of flaps: 1
Peer ID: 10.245.14.182 Local ID: 10.245.14.176 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 all end-of-rib markers sent: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  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 is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
```

```

RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 28   Sent 28   Checked 28
Input messages: Total 2      Updates 0      Refreshes 0      Octets 86
Output messages: Total 13    Updates 10     Refreshes 0      Octets 1073
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

user@PE1> show route instance detail
master:
  Router ID: 10.245.14.176
  Type: forwarding      State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
  BGP-INET:
    Router ID: 10.96.103.1
    Type: vrf            State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.245.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0      : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
  L2VPN:
    Router ID: 0.0.0.0
    Type: l2vpn          State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.512
    Route-distinguisher: 10.245.14.176:512
    Vrf-import: [ L2VPN-import ]
    Vrf-export: [ L2VPN-export ]
    Tables:

```

```

    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.96.101.1
  Type: vrf                State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.245.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0           : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.96.102.1
  Type: vrf                State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.245.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0            : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.96.100.1
  Type: vrf                State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.245.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0         : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
__juniper_private1__:
  Router ID: 0.0.0.0
  Type: forwarding        State: Active

user@PE1> show route protocol l2vpn
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
Restart Complete
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
OSPF.inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
mpls.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
800003                *[L2VPN/7] 00:06:00

```



```

> via t3-0/0/0.512, Pop      Offset: 4
t3-0/0/0.512      *[L2VPN/7] 00:06:00
> via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4
bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Complete
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.245.14.176:512:512:611/96
      *[L2VPN/7] 00:06:01
      Discard

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

Router PE1 Status During a Restart Before you can verify that graceful restart is working, you must simulate a router restart. To cause the routing process to refresh and simulate a restart, use the **restart routing** operational mode command:

```

user@PE1> restart routing
Routing protocol daemon started, pid 3558

```

The following sample output is captured during the router restart:

```

user@PE1> show bgp neighbor
Peer: 10.96.103.2      AS 65100 Local: 10.96.103.1      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: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.245.14.182+179 AS 69   Local: 10.245.14.176+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: 10.245.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.14.182   Local ID: 10.245.14.176   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 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
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

```

user@PE1> show route instance detail
master:
  Router ID: 10.245.14.176
  Type: forwarding          State: Active
  Restart State: Pending    Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.96.103.1
  Type: vrf                 State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.245.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn               State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.245.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
OSPF:
  Router ID: 10.96.101.1
  Type: vrf                 State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.245.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.96.102.1
  Type: vrf                 State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:

```

```

t3-0/0/0.102
Route-distinguisher: 10.245.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0          : 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN
STATIC:
Router ID: 10.96.100.1
Type: vrf              State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.245.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0       : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN
__juniper_private1__:
Router ID: 0.0.0.0
Type: forwarding      State: Active

```

user@PE1> show route instance summary

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding	inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		l3vpn.0	0/0/0
		inet6.0	2/0/0
		l2vpn.0	0/0/0
		l2circuit.0	0/0/0
BGP-INET	vrf	BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
L2VPN	l2vpn	L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.l2vpn.0	2/0/0
OSPF	vrf	OSPF.inet.0	7/0/0
		OSPF.iso.0	0/0/0
		OSPF.inet6.0	0/0/0
RIP	vrf	RIP.inet.0	6/0/0
		RIP.iso.0	0/0/0
		RIP.inet6.0	0/0/0
STATIC	vrf	STATIC.inet.0	4/0/0
		STATIC.iso.0	0/0/0
		STATIC.inet6.0	0/0/0
__juniper_private1__	forwarding	__juniper_priva.inet.0	0/0/0
		__juniper_privat.iso.0	0/0/0
		__juniper_priv.inet6.0	0/0/0

user@PE1> show route protocol l2vpn

```

inet.0: 16 destinations, 17 routes (15 active, 1 holddown, 1 hidden)
Restart Pending: OSPF LDP

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: OSPF LDP

BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
Restart Pending: VPN

OSPF.inet.0: 7 destinations, 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN

RIP.inet.0: 6 destinations, 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN

STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
Restart Pending: LDP VPN
+ = Active Route, - = Last Active, * = Both

800001          *[L2VPN/7] 00:00:13
                 > via t3-0/0/0.512, Pop          Offset: 4
t3-0/0/0.512    *[L2VPN/7] 00:00:13
                 > via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: VPN L2VPN
+ = Active Route, - = Last Active, * = Both

10.245.14.176:512:512:611/96
                 *[L2VPN/7] 00:00:13
                 Discard
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

```

Related Documentation

- [Enabling Graceful Restart on page 79](#)
- [Configuring Routing Protocols Graceful Restart on page 80](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 86](#)
- [Configuring VPN Graceful Restart on page 88](#)
- [Configuring Logical System Graceful Restart on page 89](#)
- [Verifying Graceful Restart Operation on page 234](#)

Example: Managing Helper Modes for OSPF Graceful Restart

- [Requirements on page 116](#)
- [Overview on page 116](#)
- [Configuration on page 116](#)
- [Verification on page 117](#)

Requirements

M Series or T Series routers running Junos OS Release 11.4 or later and EX Series switches.

Overview

Junos OS Release 11.4 extends OSPF graceful restart support to include restart signaling-based helper mode. Both standard (RFC 3623-based) and restart signaling-based helper modes are enabled by default, irrespective of the graceful-restart configuration status on the routing device.

Junos OS, however, enables you to choose between the helper modes with the **helper-disable <standard | restart-signaling | both>** statement.

Configuration

Both standard and restart signaling-based helper modes are enabled by default, irrespective of the graceful-restart configuration status on the routing device. Junos OS allows you to disable or enable the helper modes based on your requirements.

To configure the helper mode options for graceful restart:

1. To enable graceful restart, add the **graceful-restart** statement at the **[edit routing-options]** hierarchy level.

```
[edit routing-options]
user@host# set graceful-restart
```

The helper modes, both standard and restart signaling-based, are enabled by default.

2. To disable one or both of the helper modes, add the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level.

- To disable both standard and restart signaling-based helper modes:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable both
```

- To disable only the restart signaling-based helper mode:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable restart-signaling
```

- To disable only the standard helper mode:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable standard
```



NOTE: You must commit the configuration before the change takes effect.

The last committed statement always takes precedence over the previous one.

3. To enable one or both of the helper modes when the helper modes are disabled, delete the **helper-disable <both | restart-signaling | standard>** statement from the **[edit protocols ospf graceful-restart]** hierarchy level.

- To enable both standard and restart signaling-based helper modes:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable
```

- To enable the restart signaling-based helper mode:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable restart-signaling
```

- To enable the standard helper mode:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable standard
```



NOTE: You must commit the configuration before the change takes effect.

The last committed statement always takes precedence over the previous one.

Verification

Confirm that the configuration is working properly.

- [Verifying OSPF Graceful Restart and Helper Mode Configuration on page 117](#)

Verifying OSPF Graceful Restart and Helper Mode Configuration

Purpose Verify the OSPF graceful restart and helper mode configuration on a router.

- Action**
- Enter the `run show ospf overview` command from configuration mode.

```
user@host# run show ospf overview
```

```
~
~
~
Restart: Enabled
  Restart duration: 180 sec
  Restart grace period: 210 sec
  Graceful restart helper mode: Enabled
  Restart-signaling helper mode: Enabled
~
~
~
```

Meaning The output shows that graceful restart and both of the helper modes are enabled.

- Related Documentation**
- Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart
 - [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 234](#)
 - helper-disable (OSPF)

Configuration Statements: Graceful Restart

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 isis 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>
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 on page 79 • Configuring Routing Protocols Graceful Restart on page 80 • Configuring Graceful Restart for MPLS-Related Protocols on page 86 • Configuring VPN Graceful Restart on page 88 • Configuring Logical System Graceful Restart on page 89 • Graceful Restart Configuration Statements • Configuring Graceful Restart for QFabric Systems

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 in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.</p> <p>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.</p> <p>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.</p>
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	<ul style="list-style-type: none">• Enabling Graceful Restart on page 79• Configuring Routing Protocols Graceful Restart on page 80• Configuring Graceful Restart for MPLS-Related Protocols on page 86• Configuring VPN Graceful Restart on page 88• Configuring Logical System Graceful Restart on page 89• Graceful Restart Configuration Statements• Configuring Graceful Restart for QFabric Systems

helper-disable (Multiple Protocols)

Syntax	helper-disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (isis ldp ospf ospf3 rsvp) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart], [edit protocols (isis ldp ospf ospf3 rsvp) graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Disable helper mode for graceful restart. When helper mode is disabled, a router or switch cannot help a neighboring router that is attempting to restart.
Default	Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.
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 80 • Configuring Graceful Restart for MPLS-Related Protocols on page 86

maximum-helper-recovery-time

Syntax	maximum-helper-recovery-time <i>seconds</i> ;
Hierarchy Level	[edit protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the length of time the router or switch retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.
Options	<p><i>seconds</i>—Length of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.</p> <p>Range: 1 through 3600</p> <p>Default: 180</p>
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 RSVP, CCC, and TCC on page 87 • maximum-helper-restart-time (RSVP) on page 122

maximum-helper-restart-time (RSVP)

Syntax	maximum-helper-restart-time <i>seconds</i> ;
Hierarchy Level	[edit protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	Specify the length of time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. This value is applied to all RSVP neighbor routers and should be based on the time that the slowest RSVP neighbor requires for restart.
Options	seconds —The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. Range: 1 through 1800 Default: 60
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 RSVP, CCC, and TCC on page 87• maximum-helper-recovery-time on page 121

maximum-neighbor-reconnect-time

Syntax	maximum-neighbor-reconnect-time <i>seconds</i> ;
Hierarchy Level	[edit protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Specify the maximum length of time allowed to reestablish connection from a restarting neighbor.
Options	seconds —Maximum time allowed for reconnection. Range: 30 through 300
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 LDP on page 87

maximum-neighbor-recovery-time

Syntax	<code>maximum-neighbor-recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement changed from maximum-recovery-time to maximum-neighbor-recovery-time in Junos OS Release 9.1.
Description	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
Options	seconds —Configure the maximum recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 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 LDP Graceful Restart • Configuring Graceful Restart Options for LDP on page 87 • no-strict-lsa-checking on page 124 • recovery-time on page 127

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.
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 83• Configuring Graceful Restart for QFabric Systems• maximum-neighbor-recovery-time on page 123• recovery-time on page 127

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>
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 83 • Configuring Graceful Restart for QFabric Systems • restart-duration on page 128

reconnect-time

Syntax	<code>reconnect-time <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
Options	<i>seconds</i> —Time required for reconnection. Range: 30 through 300 Default: 60 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 LDP Graceful Restart on LDP Configuration GuideConfiguring Graceful Restart Options for LDP on page 87

recovery-time

Syntax	<code>recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the length of time a router or switch waits for Label Distribution Protocol (LDP) neighbors to assist it with a graceful restart.
Options	seconds —Time the router waits for LDP to restart gracefully. Range: 120 through 1800 Default: 160
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 Graceful Restart Options for LDP on page 87 • maximum-neighbor-recovery-time on page 123 • no-strict-lsa-checking on page 124

restart-duration

Syntax	<code>restart-duration <i>seconds</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (isis ospf ospf3 pim) graceful-restart],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</code> <code>[edit protocols (esis isis ospf ospf3 pim) graceful-restart],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</code> <code>[edit routing-options graceful-restart]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
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	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Enabling Graceful Restart on page 79](#)
 - [Configuring Routing Protocols Graceful Restart on page 80](#)
 - [Configuring Graceful Restart for MPLS-Related Protocols on page 86](#)
 - [Configuring VPN Graceful Restart on page 88](#)
 - [Configuring Graceful Restart for VPNs](#)
 - [Configuring Logical System Graceful Restart on page 89](#)
 - [Graceful Restart Configuration Statements](#)
 - [Configuring Graceful Restart for QFabric Systems](#)

restart-time (BGP Graceful Restart)

Syntax	<code>restart-time <i>seconds</i>;</code>
Hierarchy Level	<p>[edit protocols (bgp rip ripng) graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols (bgp rip ripng) graceful-restart (Enabling Globally)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]</p>
Release Information	<p>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.</p>
Description	Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.
Options	<p><i>seconds</i>—Length of time for the graceful restart period. Range: 1 through 600 seconds Default: Varies by protocol:</p> <ul style="list-style-type: none"> • BGP—120 seconds • RIP and RIPng—60 seconds
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 81 • Configuring Graceful Restart Options for RIP and RIPng on page 84 • Configuring Graceful Restart for QFabric Systems • stale-routes-time on page 130

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.
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 81• Configuring Graceful Restart for QFabric Systems• restart-time (BGP Graceful Restart) on page 129

traceoptions (Protocols)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit protocols isis], [edit protocols (ospf ospf3)]
Release Information	Statement introduced before Junos OS Release 7.4. graceful-restart flag for IS-IS and OSPF/OSPFv3 added in Junos OS Release 8.4.
Description	<p>Define tracing operations that graceful restart functionality in the router or switch.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The nonstop active routing tracing option is:</p> <ul style="list-style-type: none"> • graceful-restart—Tracing operations for nonstop active routing <p>no-world-readable—Restrict users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues</p>

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 system

Default: 128 KB

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	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Tracking Graceful Restart Events on page 86

Configuration: NSB

- [Configuring Nonstop Bridging on page 133](#)
- [Resetting Local Statistics on page 134](#)

Configuring Nonstop Bridging

This section includes the following topics:

- [Enabling Nonstop Bridging on page 133](#)
- [Synchronizing the Routing Engine Configuration on page 133](#)
- [Verifying Nonstop Bridging Operation on page 134](#)

Enabling Nonstop Bridging

Nonstop bridging requires you to configure graceful Routing Engine switchover (GRES). To enable graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level:

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

By default, nonstop bridging is disabled. To enable nonstop bridging, include the **nonstop-bridging** statement at the **[edit protocols layer2-control]** hierarchy level:

```
[edit protocols layer2-control]  
nonstop-bridging;
```

To disable nonstop active routing, remove the **nonstop-bridging** statement from the **[edit protocols layer2-control]** hierarchy level.

Synchronizing the Routing Engine Configuration

When you configure nonstop bridging, you must also include the **commit synchronize** statement at the **[edit system]** hierarchy level so that, by default, when you issue the **commit** command, the configuration changes are synchronized on both Routing Engines. If you issue the **commit synchronize** command at the **[edit]** hierarchy level on the backup Routing Engine, the Junos OS displays a warning and commits the candidate configuration.



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

When you configure nonstop bridging, 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 Nonstop Bridging Operation

When you enable nonstop bridging, you can issue Layer 2 Control Protocol-related operational mode commands on the backup Routing Engine. However, the output of the commands might not match the output of the same commands issued on the master Routing Engine.

Related Documentation

- [Nonstop Bridging Concepts on page 11](#)
- [Nonstop Bridging System Requirements on page 13](#)
- [nonstop-bridging on page 136](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)

Resetting Local Statistics

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

Related Documentation

- [Configuring Nonstop Active Routing on page 137](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 139](#)

Configuration Statements: NSB

- [\[edit protocols layer2-control\] Hierarchy Level](#) on page 135

[\[edit protocols layer2-control\] Hierarchy Level](#)

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

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

Related Documentation

- Notational Conventions Used in Junos OS Configuration Hierarchies
- [\[edit protocols\] Hierarchy Level](#)

nonstop-bridging

Syntax	nonstop-bridging;
Hierarchy Level	[edit protocols layer2-control]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Synchronizing the Routing Engine Configuration on page 138• Configuring Nonstop Bridging on page 133• Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)

Configuration: NSR

- [Configuring Nonstop Active Routing on page 137](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 139](#)
- [Example: Configuring Nonstop Active Routing on page 143](#)

Configuring Nonstop Active Routing

This section includes the following topics:

- [Enabling Nonstop Active Routing on page 137](#)
- [Synchronizing the Routing Engine Configuration on page 138](#)
- [Verifying Nonstop Active Routing Operation on page 138](#)

Enabling Nonstop Active Routing

Nonstop active routing (NSR) requires you to configure graceful Routing Engine switchover (GRES). To enable graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level:

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

By default, nonstop active routing is disabled. To enable nonstop active routing, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level:

```
[edit routing-options]  
nonstop-routing;
```

To disable nonstop active routing, remove the **nonstop-routing** statement from the **[edit routing-options]** hierarchy level.



NOTE: When you enable nonstop active routing, you cannot enable automatic route distinguishers for multicast VPN routing instances. Automatic route distinguishers are enabled by configuring the **route-distinguisher-id** statement at the **[edit routing-instances *instance-name*]** hierarchy level; for more information, see the Junos OS VPNs Configuration Guide.

To enable the routing platform to switch over to the backup Routing Engine when the routing protocol process (rpd) fails rapidly three times in succession, include the

other-routing-engine statement at the **[edit system processes routing failover]** hierarchy level.

For more information about the **other-routing-engine** statement, see the Junos OS System Basics Configuration Guide.

Synchronizing the Routing Engine Configuration

When you configure nonstop active routing, you must also include the **commit synchronize** statement at the **[edit system]** hierarchy level so that configuration changes are synchronized on both Routing Engines:

```
[edit system]
commit synchronize;
```

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

If you configure the **commit synchronize** statement at the **[edit system]** hierarchy level and issue a commit in the master Routing Engine, the master configuration is automatically synchronized with the backup.

However, if the backup Routing Engine is down when you issue the commit, the Junos OS displays a warning and commits the candidate configuration in the master Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the master.



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

When you configure nonstop active routing, 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 Nonstop Active Routing Operation

To see whether or not nonstop active routing is enabled, issue the **show task replication** command. For BGP nonstop active routing, you can also issue the **show bgp replication** command.

For more information about these commands, see the Junos OS Operational Mode Commands and Junos OS Operational Mode Commands, respectively.

When you enable nonstop active routing or graceful Routing Engine switchover and issue routing-related operational mode commands on the backup Routing Engine (such as **show route**, **show bgp neighbor**, **show ospf database**, and so on), the output might not match the output of the same commands issued on the master Routing Engine. For example, it is normal for the routing table on the backup Routing Engine to contain persistent phantom routes that are not present in the routing table on the master Routing Engine.

To display BFD state replication status, issue the **show bfd session** command. The **replicated** flag appears in the output for this command when a BFD session has been replicated to the backup Routing Engine. For more information, see the Junos OS Operational Mode Commands.

Related Documentation

- [Nonstop Active Routing Concepts on page 15](#)
- [Nonstop Active Routing System Requirements on page 18](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 139](#)
- [Resetting Local Statistics on page 134](#)
- [Example: Configuring Nonstop Active Routing on page 143](#)
- [nonstop-routing on page 149](#)

Tracing Nonstop Active Routing Synchronization Events

To track the progress of nonstop active routing synchronization between Routing Engines, you can configure nonstop active routing trace options flags for each supported protocol and for BFD sessions and record these operations to a log file.

To configure nonstop active routing trace options for supported routing protocols, include the **nsr-synchronization** statement at the **[edit protocols *protocol-name* traceoptions flag]** hierarchy level and optionally specify one or more of the **detail**, **disable**, **receive**, and **send** options:

```
[edit protocols]
  bgp {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  isis {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  ldp {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  mpls {
    traceoptions {
      flag nsr-synchronization;
      flag nsr-synchronization-detail;
    }
  }
  msdp {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
}
```

```
(ospf | ospf3) {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
rip {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
ripng {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
pim {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
```

To configure nonstop active routing trace options for BFD sessions, include the **nsr-synchronization** and **nsr-packet** statements at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```
[edit protocols]
bfd {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-packet;
  }
}
```

To trace the Layer 2 VPN signaling state replicated from routes advertised by BGP, include the **nsr-synchronization** statement at the **[edit routing-options traceoptions flag]** hierarchy level. This flag also traces the label and logical interface association that VPLS receives from the kernel replication state.

```
[edit routing-options]
traceoptions {
  flag nsr-synchronization;
}
```

Related Documentation

- [Configuring Nonstop Active Routing on page 137](#)
- [Configuring Nonstop Active Routing on EX Series Switches \(CLI Procedure\)](#)
- [Example: Configuring Nonstop Active Routing on page 143](#)
- [Example: Configuring Nonstop Active Routing on EX Series Switches](#)

traceoptions (Routing Options)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>Values:</p> <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(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</p>

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

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

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

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

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

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

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Tracing Global Routing Protocol Operations • Tracing Nonstop Active Routing Synchronization Events on page 139

Example: Configuring Nonstop Active Routing

The following example enables graceful Routing Engine switchover, nonstop active routing, and nonstop active routing trace options for BGP, IS-IS, and OSPF.

```
[edit]
system commit {
  synchronize;
}
chassis {
  redundancy {
    graceful-switchover; # This enables graceful Routing Engine switchover on
                        # the routing platform.
  }
}
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.0.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.2.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.3.1.1/30;
      }
    }
  }
}
```

```
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.2.1/32;
        }
        family iso {
            address 49.0004.1921.6800.2001.00;
        }
    }
}
}
routing-options {
    nonstop-routing; # This enables nonstop active routing on the routing platform.
    router-id 192.168.2.1;
    autonomous-system 65432;
}
protocols {
    bgp {
        traceoptions {
            flag nsr-synchronization detail; # This logs nonstop active routing
            # events for BGP.
        }
        local-address 192.168.2.1;
        group external-group {
            type external;
            export BGP_export;
            neighbor 192.168.1.1 {
                family inet {
                    unicast;
                }
                peer-as 65103;
            }
        }
        group internal-group {
            type internal;
            neighbor 192.168.10.1;
            neighbor 192.168.11.1;
            neighbor 192.168.12.1;
        }
    }
}
isis {
    traceoptions {
        flag nsr-synchronization detail; # This logs nonstop active routing events
        # for IS-IS.
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        passive;
    }
}
```

```
ospf {
  traceoptions {
    flag nsr-synchronization detail; # This logs nonstop active routing events
    # for OSPF.
  }
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface lo0.0 {
      passive;
    }
  }
}
}
policy-options {
  policy-statement BGP_export {
    term direct {
      from {
        protocol direct;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
}
```

- Related Documentation**
- [Configuring Nonstop Active Routing on page 137](#)
 - [Tracing Nonstop Active Routing Synchronization Events on page 139](#)

Configuration Statements: NSR

- [\[edit protocols layer2-control\] Hierarchy Level](#) on page 147

[\[edit protocols layer2-control\] Hierarchy Level](#)


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

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

Related Documentation

- Notational Conventions Used in Junos OS Configuration Hierarchies
- [\[edit protocols\] Hierarchy Level](#)

commit synchronize

Syntax	commit synchronize;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 10.4 for EX Series switches.
Description	<p>For devices with multiple Routing Engines only. Configure the commit command to automatically result in a commit synchronize action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the commit command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.</p>
	<div> NOTE: When you configure nonstop active routing (NSR), you must include the commit synchronize statement. Otherwise, the commit operation fails.</div>
	<p>On the TX Matrix router, synchronization only occurs between the Routing Engines within the same chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the T640 routers. That is, the master Routing Engine on the TX Matrix router distributes the configuration to the master Routing Engine on each T640 router. Likewise, the backup Routing Engine on the TX Matrix router distributes the configuration to the backup Routing Engine on each T640 router.</p> <p>In EX Series Virtual Chassis configurations:</p> <ul style="list-style-type: none">• On EX4200 switches in Virtual Chassis, synchronization occurs between the switch in the master role and the switch in the backup role.• On EX8200 switches in a Virtual Chassis, synchronization occurs only between the master and backup XRE200 External Routing Engines.
Options	<p>and-quit—(Optional) (EX Series only) Quit configuration mode if the commit synchronization succeeds.</p> <p>comment—(Optional) (EX Series only) Write a message to the commit log.</p> <p>and-force—(Optional) (EX Series only) Force a commit synchronization on the other Routing Engine (ignore warnings).</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Synchronizing the Routing Engine Configuration on page 138

- Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically

nonstop-routing

Syntax	nonstop-routing;
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-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 10.4 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.
Default	disabled
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Nonstop Active Routing on page 137 • Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure)

traceoptions (Routing Options)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>Values:</p> <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(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</p>

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

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

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

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

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

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

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

Related Documentation	<ul style="list-style-type: none">• Example: Tracing Global Routing Protocol Operations• Tracing Nonstop Active Routing Synchronization Events on page 139
------------------------------	---

CHAPTER 15

Configuration: Unified ISSU

- [Best Practices on page 153](#)
- [Before You Begin on page 154](#)
- [Performing a Unified ISSU on page 157](#)
- [Verifying a Unified ISSU on page 169](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 170](#)

Best Practices

When you are planning to perform a unified in-service software upgrade (unified ISSU), choose a time when your network is as stable as possible. As with a normal upgrade, Telnet sessions, SNMP, and CLI access are briefly interrupted. In addition, the following restrictions apply:

- The master Routing Engine and backup Routing Engine must be running the same software version before you can perform a unified ISSU.
- During a unified ISSU, you cannot bring any PICs online or offline.
- Unicast RPF-related statistics are not saved across a unified ISSU, and the unicast RPF counters are reset to zero during a unified ISSU.

Related Documentation

- [Before You Begin on page 154](#)
- [Performing a Unified ISSU on page 157](#)
- [Verifying a Unified ISSU on page 169](#)
- [Troubleshooting Unified ISSU Problems on page 247](#)

Before You Begin

Before you begin a unified ISSU, complete the tasks in the following sections:

1. [Verify That the Master and Backup Routing Engines Are Running the Same Software Version on page 155](#)
2. [Back Up the Router Software on page 155](#)
3. [Verify That Graceful Routing Engine Switchover and Nonstop Active Routing Are Configured on page 156](#)

Verify That the Master and Backup Routing Engines Are Running the Same Software Version

To verify that both Routing Engines are running the same version of software, issue the following command:

```
{master}
user@host> show version invoke-on all-routing-engines
re0:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071210.0]
JUNOS Base OS Software Suite [9.0-20071210.0]
JUNOS Kernel Software Suite [9.0-20071210.0]
JUNOS Crypto Software Suite [9.0-20071210.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071210.0]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071210.0]
JUNOS Online Documentation [9.0-20071210.0]
JUNOS Routing Software Suite [9.0-20071210.0]
re1:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071210.0]
JUNOS Base OS Software Suite [9.0-20071210.0]
JUNOS Kernel Software Suite [9.0-20071210.0]
JUNOS Crypto Software Suite [9.0-20071210.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071210.0]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071210.0]
JUNOS Online Documentation [9.0-20071210.0]
JUNOS Routing Software Suite [9.0-20071210.0]
```

If both Routing Engines are not running the same software version, issue the **request system software add** command on the desired Routing Engine so that the software version is the same. For more information, see the Installation and Upgrade Guide.

Back Up the Router Software

As a preventive measure in case any problems occur during an upgrade, issue the **request system snapshot** command on *each* Routing Engine to back up the system software to the router's hard disk. The following is an example of issuing the command on the master Routing Engine:

```
{master}
user@host> request system snapshot
Verifying compatibility of destination media partitions...
Running newfs (220MB) on hard-disk media / partition (ad1s1a)...
Running newfs (24MB) on hard-disk media /config partition (ad1s1e)...
Copying '/dev/ad0s1a' to '/dev/ad1s1a' .. (this may take a few minutes)
Copying '/dev/ad0s1e' to '/dev/ad1s1e' .. (this may take a few minutes)
The following filesystems were archived: / /config
```



NOTE: The root file system is backed up to `/altroot`, and `/config` is backed up to `/altconfig`. After you issue the `request system snapshot` command, the router's flash and hard disks are identical. You can return to the previous version of the software only by booting the router from removable media. For more information about the `request system snapshot` command, see the Junos OS System Basics Configuration Guide.

Verify That Graceful Routing Engine Switchover and Nonstop Active Routing Are Configured

Before you begin a unified ISSU, ensure that graceful Routing Engine switchover and nonstop active routing are configured on your router.

1. To verify graceful Routing Engine switchover is configured, on the backup Routing Engine (**re1**) issue the `show system switchover` command. The output should be similar to the following example. The **Graceful switchover** field state must be **On**.

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

2. To verify nonstop active routing is configured, on the master Routing Engine (**re0**) issue the `show task replication` command. The output should be similar to the following example.

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

Protocol           Synchronization Status
OSPF                Complete
IS-IS               Complete
```

If graceful Routing Engine switchover and nonstop active routing are not configured, complete the following steps:

1. On the master Routing Engine (**re0**), enable graceful Routing Engine switchover. Include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.
2. On the master Routing Engine, enable nonstop active routing. Include the **commit synchronize** statement at the **[edit system]** hierarchy level and the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level.
3. On the master Router Engine, issue the **commit** command.

The system provides the following confirmation that the master and backup Routing Engines are synchronized:

```
re0:
configuration check succeeds
```

```

rel:
commit complete
re0:
commit complete

```

Related Documentation

- [Unified ISSU Concepts on page 37](#)
- [Unified ISSU Process on the TX Matrix Router](#)
- [Unified ISSU System Requirements on page 42](#)
- [Best Practices on page 153](#)
- [Performing a Unified ISSU on page 157](#)
- [Verifying a Unified ISSU on page 169](#)
- [Troubleshooting Unified ISSU Problems on page 247](#)

Performing a Unified ISSU

You can perform a unified ISSU in one of three ways:

1. [Upgrading and Rebooting Both Routing Engines Automatically on page 157](#)
2. [Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually on page 161](#)
3. [Upgrading and Rebooting Only One Routing Engine on page 166](#)

Upgrading and Rebooting Both Routing Engines Automatically

When you issue the **request system software in-service-upgrade** command with the **reboot** option, the system automatically upgrades both Routing Engines to the newer software and reboots both Routing Engines. This option enables you to complete the unified ISSU with a single command.

To perform a unified ISSU using the **request system software in-service-upgrade package-name reboot** command, complete the following steps:

1. Download the software package from the Juniper Networks Support website, <http://www.juniper.net/support/>. Choose the Canada and U.S., Worldwide, or Junos-FIPS edition. Place the package on a local server. To download the package, you must have a service contract and an access account. If you do not have an access account, complete the registration form at the Juniper Networks website: <https://www.juniper.net/registration/Register.jsp>.
2. Copy the package to the router. We recommend that you copy it to the **/var/tmp** directory, which is a large file system on the hard disk.


```
user@host>file copy ftp://username:prompt@ftp.hostname.net/filename/var/tmp/filename
```
3. To verify the current software version running on both Routing Engines, on the master Routing Engine issue the **show version invoke-on all-routing-engines** command. The following example shows that both Routing Engines are running an image of Junos OS, Release 9.0, that was built on December 11, 2007:

```
{backup}
```

```
user@host> show version invoke-on all-routing-engines
re0:
```

```
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite 9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.2]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]
```

```
re1:
```

```
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite [9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.20]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]
```

4. On the master Routing Engine, issue the **request system software in-service-upgrade package-name reboot** command. The following example upgrades the current version to an image of Junos OS, Release 9.0, that was built on January 14, 2008:

```
{master}
```

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz reboot
ISSU: Validating Image
PIC 0/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080114.2
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0-20080114.2-domestic.tgz
Using jbundle-9.0-20080114.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpfe-9.0-20080114.2.tgz
Using jdocs-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0-20080114.2.tgz
```



```

Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING: This package will load JUNOS 9.0-20080114.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz
...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 6         Online (ISSU)
  FPC 7         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/paKEuy' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...

```

Verified manifest signed by PackageProduction_9_0_0

```
WARNING: This package will load JUNOS 9.0-20080114.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.
```

Saving the config files ...

NOTICE: uncommitted changes have been saved in

/var/db/config/juniper.conf.pre-install

Installing the bootstrap installer ...

```
WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.
```

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz

...

cp: /var/tmp/paKEuy is a directory (not copied).

Saving state for rollback ...

ISSU: Old Master Upgrade Done

ISSU: IDLE

Shutdown NOW!

Reboot consistency check bypassed - jinstall 9.0-20080114.2 will complete
installation upon reboot

[pid 30227]

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

System going down IMMEDIATELY

Connection to host closed.

When the new backup (old master) Routing Engine is rebooted, you are logged out from the router.

5. After waiting a few minutes, log in to the router again. You are logged in to the new backup Routing Engine (**re0**). To verify that both Routing Engines have been upgraded, issue the following command:

```
{backup}
```

```
user@host> show version invoke-on all-routing-engines
```

```
re0:
```

```
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20080114.2]
JUNOS Base OS Software Suite 9.0-20080114.2]
JUNOS Kernel Software Suite [9.0-20080114.2]
JUNOS Crypto Software Suite [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20080114.2]
JUNOS Online Documentation [9.0-20080114.2]
JUNOS Routing Software Suite [9.0-20080114.2]
```

```
rel:
```

```
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0-20080114.2]
JUNOS Base OS Software Suite [9.0-20080114.2]
JUNOS Kernel Software Suite [9.0-20080114.2]
JUNOS Crypto Software Suite [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20080114.2]
JUNOS Online Documentation [9.0-20080114.2]
JUNOS Routing Software Suite [9.0-20080114.2]
```

6. To make **re0** the master Routing Engine, issue the following command:

```
{backup}

user@host> request chassis routing-engine master acquire
Attempt to become the master routing engine ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.

{master}
user@host>
```

7. Issue the **request system snapshot** command on *each* Routing Engine to back up the system software to the router's hard disk.



NOTE: The root file system is backed up to `/altroot`, and `/config` is backed up to `/altconfig`. After you issue the **request system snapshot** command, the router's flash and hard disks are identical. You can return to the previous version of the software only by booting the router from removable media.

Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually

When you issue the **request system software in-service-upgrade** command without any options, the system upgrades and reboots the new master Routing Engine to the newer software. The new software is placed on the new backup (old master) Routing Engine; however, to complete the upgrade, you must issue the **request system reboot** command on the new backup Routing Engine.

To perform a unified ISSU using the **request system software in-service-upgrade package-name** command without any options, complete the following steps:

1. Download the software package from the Juniper Networks Support website, <http://www.juniper.net/support/>. Choose the Canada and U.S., Worldwide, or Junos-FIPS edition. Place the package on a local server. To download the package, you must have a service contract and an access account. If you do not have an access account, complete the registration form at the Juniper Networks website: <https://www.juniper.net/registration/Register.jsp>.

2. Copy the package to the router. We recommend that you copy it to the **/var/tmp** directory, which is a large file system on the hard disk.

```
user@host>file copy ftp://username:prompt@ftp.hostname.net/filename/var/tmp/filename
```

3. To verify the current software version running on both Routing Engines, on the master Routing Engine, issue the **show version invoke-on all-routing-engines** command. The following example shows that both Routing Engines are running Junos OS Release 9.0R1:

```
{master}
```

```
user@host> show version invoke-on all-routing-engines
re0:
```

```
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
JUNOS Online Documentation [9.0R1]
JUNOS Routing Software Suite [9.0R1]
```

```
re1:
```

```
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
JUNOS Online Documentation [9.0R1]
JUNOS Routing Software Suite [9.0R1]
```

4. On the master Routing Engine, issue the **request system software in-service-upgrade package-name** command. The following example upgrades the current version to Junos OS Release 9.0R1.2:

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0R1.2-domestic-signed.tgz
ISSU: Validating Image
FPC 4 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

ISSU: Preparing Backup RE
```

```

Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080117.0
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0R1.2-domestic-signed.tgz
Verified jinstall-9.0R1.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0R1.2-domestic.tgz
Using jbundle-9.0R1.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpf-9.0R1.2.tgz
Using jdocs-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0R1.2-domestic-signed.tgz' ...
Verified jinstall-9.0R1.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING: This package will load JUNOS 9.0R1.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0R1.2-domestic-signed.tgz ...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done

```

```

ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status           Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Offline          Offlined by cli command
  FPC 5          Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/paeBi5' ...
Verified jinstall-9.0R1.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING: This package will load JUNOS 9.0R1.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0R1.2-domestic-signed.tgz ...
cp: /var/tmp/paeBi5 is a directory (not copied).
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE

```

5. Issue the **show version invoke-on all-routing-engines** command to verify that the new backup (old master) Routing Engine (**re0**), is still running the previous software image, while the new master Routing Engine (**re1**) is running the new software image:

```

{backup}

user@host> show version
re0:
-----
Hostname: user
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]

```

```

JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
JUNOS Online Documentation [9.0R1]
JUNOS Routing Software Suite [9.0R1]
labpkg [7.0]
JUNOS Installation Software [9.0R1.2]

```

```
re1:
```

```

-----
Hostname: user1
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
JUNOS Online Documentation [9.0R1.2]
JUNOS Routing Software Suite [9.0R1.2]

```

6. At this point, if you choose not to install the newer software version on the new backup Routing Engine (**re1**), you can issue the **request system software delete jinstall** command on it. Otherwise, to complete the upgrade, go to the next step.
7. Reboot the new backup Routing Engine (**re0**) by issuing the **request system reboot** command:

```
{backup}
```

```

user@host> request system reboot
Reboot the system ? [yes,no] (no) yes

```

```

Shutdown NOW!
Reboot consistency check bypassed - jinstall 9.0R1.2 will complete installation
upon reboot
[pid 6170]

```

```
{backup}
```

```

user@host>
System going down IMMEDIATELY

```

```

Connection to host closed by remote host.
Connection to host closed.

```

If you are not on the console port, you are disconnected from the router session.

8. After waiting a few minutes, log in to the router again. You are logged in to the new backup Routing Engine (**re0**). To verify that both Routing Engines have been upgraded, issue the following command:

```
{backup}
```

```

user@host> show version invoke-on all-routing-engines
re0:

```

```

-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]

```

```
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
JUNOS Online Documentation [9.0R1.2]
JUNOS Routing Software Suite [9.0R1.2]
```

```
re1:
```

```
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
JUNOS Online Documentation [9.0R1.2]
JUNOS Routing Software Suite [9.0R1.2]
```

9. To make **re0** the master Routing Engine, issue the following command:

```
{backup}
```

```
user@host> request chassis routing-engine master acquire
Attempt to become the master routing engine ? [yes,no] (no) yes
```

```
Resolving mastership...
Complete. The local routing engine becomes the master.
```

```
{master}
user@host>
```

10. Issue the **request system snapshot** command on *each* Routing Engine to back up the system software to the router's hard disk.



NOTE: The root file system is backed up to `/altroot`, and `/config` is backed up to `/altconfig`. After you issue the **request system snapshot** command, the router's flash and hard disks are identical. You can return to the previous version of the software only by booting the router from removable media.

Upgrading and Rebooting Only One Routing Engine

When you issue the **request system software in-service-upgrade** command with the **no-old-master-upgrade** option, the system upgrades and reboots only the new master Routing Engine. To upgrade the new backup (former master) Routing Engine, you must issue the **request system software add** command.

To perform a unified ISSU using the **request system software in-service-upgrade package-name no-old-master-upgrade** commands, complete the following steps:

1. Download the software package from the Juniper Networks Support website, <http://www.juniper.net/support/>. Choose the Canada and U.S., Worldwide, or Junos-FIPS edition. Place the package on a local server. To download the package, you must have a service contract and an access account. If you do not have an access account, complete the registration form at the Juniper Networks website: <https://www.juniper.net/registration/Register.jsp>.

2. Copy the package to the router. We recommend that you copy it to the **/var/tmp** directory, which is a large file system on the hard disk.

```
user@host>file copy ftp://username:prompt@ftp.hostname.net/filename/var/tmp/filename
```

3. To verify the current software version running on both Routing Engines, on the master Routing Engine issue the **show version invoke-on all-routing-engines** command. The following example shows that both Routing Engines are running an image of Junos OS Release 9.0 that was built on December 11, 2007:

```
{backup}
```

```
user@host> show version invoke-on all-routing-engines
re0:
```

```
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite 9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.2]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]
```

```
re1:
```

```
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite [9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.20]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]
```

4. On the master Routing Engine, issue the **request system software in-service-upgrade package-name no-old-master-upgrade** command. The following example upgrades the current version to an image of Junos OS Release 9.0 that was built on January 16, 2008:

```
{master}
```

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz no-old-master-upgrade
ISSU: Validating Image
ISSU: Preparing Backup RE
```

```
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080116.2
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz
Verified jinstall-9.0-20080116.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0-20080116.2-domestic.tgz
Using jbundle-9.0-20080116.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpfe-9.0-20080116.2.tgz
Using jdocs-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz' ...
Verified jinstall-9.0-20080116.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING:   This package will load JUNOS 9.0-20080116.2 software.
WARNING:   It will save JUNOS configuration files, and SSH keys
WARNING:   (if configured), but erase all other files and information
WARNING:   stored on this machine. It will attempt to preserve dumps
WARNING:   and log files, but this can not be guaranteed. This is the
WARNING:   pre-installation stage and all the software is loaded when
WARNING:   you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:   A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:   'request system reboot' command when software installation is
WARNING:   complete. To abort the installation, do not reboot your system,
WARNING:   instead use the 'request system software delete jinstall'
WARNING:   command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080116.2-domestic-signed.tgz
...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
```

```

ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 3         Online (ISSU)
  FPC 5         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE

{backup}
user@host>

```

5. You are now logged in to the new backup (old master Routing Engine). If you want to install the new software version on the new backup Routing Engine, issue the **request system software add /var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz** command.

Related Documentation

- [Unified ISSU System Requirements on page 42](#)
- [Best Practices on page 153](#)
- [Before You Begin on page 154](#)
- [Verifying a Unified ISSU on page 169](#)
- [Troubleshooting Unified ISSU Problems on page 247](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 170](#)

Verifying a Unified ISSU

To verify the status of FPCs and their corresponding PICs after the most recent unified ISSU, issue the **show chassis in-service-upgrade** command on the master Routing Engine:

```

user@host> show chassis in-service-upgrade
  Item          Status          Reason
  FPC 0         Online
  FPC 1         Online
  FPC 2         Online
  PIC 0         Online
  PIC 1         Online
  FPC 3         Offline          Offlined by CLI command
  FPC 4         Online
  PIC 1         Online
  FPC 5         Online
  PIC 0         Online
  FPC 6         Online

```

PIC 3 Online
FPC 7 Online

For more information about the **show chassis in-service-upgrade** command, see the Junos OS Operational Mode Commands.

**Related
Documentation**

- [Performing a Unified ISSU on page 157](#)
- [Troubleshooting Unified ISSU Problems on page 247](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 170](#)

Managing and Tracing BFD Sessions During Unified ISSU Procedures

Bidirectional Forwarding Detection (BFD) sessions temporarily increase their detection and transmission timers during unified ISSU procedures. After the upgrade, these timers revert to the values in use before the unified ISSU started. The BFD process replicates the unified ISSU state and timer values to the backup Routing Engine for each session.

No additional configuration is necessary to enable unified ISSU for BFD. However, you can disable the BFD timer negotiation during the unified ISSU by including the **no-issu-timer-negotiation** statement at the **[edit protocols bfd]** hierarchy level:

```
[edit protocols bfd]  
no-issu-timer-negotiation;
```

If you configure this statement, the BFD timers maintain their original values during unified ISSU.



CAUTION: The sessions might flap during unified ISSU or Routing Engine switchover, depending on the detection intervals.

For more information about BFD, see the Junos OS Routing Protocols Configuration Guide.

To configure unified ISSU trace options for BFD sessions, include the **issu** statement at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```
[edit protocols]  
bfd {  
  traceoptions {  
    flag issu;  
  }  
}
```

**Related
Documentation**

- [Unified ISSU Concepts on page 37](#)
- [Unified ISSU Process on the TX Matrix Router](#)
- [Unified ISSU System Requirements on page 42](#)
- [Best Practices on page 153](#)
- [Before You Begin on page 154](#)

- [Performing a Unified ISSU on page 157](#)
- [Verifying a Unified ISSU on page 169](#)
- [Troubleshooting Unified ISSU Problems on page 247](#)

CHAPTER 16

Configuration Statements: Unified ISSU

bfd

Syntax	<pre>bfd { traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure trace options for Bidirectional Forwarding Protocol (BFD) traffic.
Default	If you do not include this statement, no BFD tracing operations are performed.
Options	<p>disable—(Optional) Disable the BFD tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the /var/log directory. We recommend that you place global routing protocol tracing output in the routing-log file.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. These are the BFD protocol tracing options:</p> <ul style="list-style-type: none">• adjacency—Trace adjacency messages.• all—Trace all options for BFD.• error—Trace all errors.• event—Trace all events.• issu—Trace in-service software upgrade (ISSU) packet activity.

- **nsr-packet**—Trace non-stop-routing (NSR) packet activity.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management (PPM).
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

match *regular-expression*—(Optional) Regular expression for lines to be logged.

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

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the trace file again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

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

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

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

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD for Static Routes

no-issu-timer-negotiation

Syntax	no-issu-timer-negotiation;
Hierarchy Level	[edit protocols bfd], [edit logical-systems <i>logical-system-name</i> protocols bfd], [edit routing-instances <i>routing-instance-name</i> protocols bfd]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Disable unified ISSU timer negotiation for Bidirectional Forwarding Detection (BFD) sessions.



CAUTION: The sessions might flap during unified ISSU or Routing Engine switchover, depending on the detection intervals.

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">• Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 170• Junos OS Routing Protocols Configuration Guide.

traceoptions (Protocols BFD)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit protocols bfd]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>issu flag for BFD added in Junos OS Release 9.1.</p>
Description	<p>Define tracing operations that track unified in-service software upgrade (ISSU) functionality in the router or switch.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. There is only one unified ISSU tracing option:</p> <ul style="list-style-type: none"> issu—Trace BFD unified ISSU operations. <p>no-world-readable—Restrict users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <i>trace-file</i> reaches this size, it is renamed <i>trace-file.0</i>. When the <i>trace-file</i> again reaches its maximum size, <i>trace-file.0</i> is renamed <i>trace-file.1</i> and <i>trace-file</i> is renamed <i>trace-file.0</i>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p>

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

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

Related Documentation	<ul style="list-style-type: none">• Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 170
------------------------------	--

CHAPTER 17

Configuration: VRRP

- [Configuring the Startup Period for VRRP Operations on page 179](#)
- [Configuring Basic VRRP Support on page 180](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 182](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 184](#)
- [Configuring a Backup Router to Preempt the Master Router on page 186](#)
- [Modifying the Preemption Hold-Time Value on page 187](#)
- [Configuring Asymmetric Hold Time for VRRP Routers on page 187](#)
- [Configuring an Interface to Accept Packets Destined for the Virtual IP Address on page 188](#)
- [Configuring a Logical Interface to Be Tracked on page 189](#)
- [Configuring a Route to Be Tracked on page 191](#)
- [Configuring Inheritance for a VRRP Group on page 192](#)
- [Configuring the Silent Period on page 193](#)
- [Configuring Passive ARP Learning for Backup VRRP Routers on page 194](#)
- [Enabling the Distributed Periodic Packet Management Process for VRRP on page 195](#)
- [Configuring VRRP to Improve Convergence Time on page 196](#)
- [Example: Configuring VRRP on page 197](#)
- [Example: Configuring VRRP for IPv6 on page 199](#)
- [Example: Configuring VRRP Route Tracking on page 200](#)
- [Tracing VRRP Operations on page 201](#)

Configuring the Startup Period for VRRP Operations

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;
```



NOTE: During the silent startup period, the `show vrrp detail` command output shows a value of 0 for Master priority, and your own IP address for Master router. These values indicate that the Master selection is not completed yet, and these values can be ignored.

**Related
Documentation**

- [Understanding VRRP on page 55](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring Basic VRRP Support on page 180](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 182](#)
- [Example: Configuring VRRP on page 197](#)

Configuring Basic VRRP Support

An interface can be a member of one or more VRRP groups. 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 ];  
}
```

You can include this statement at the following hierarchy levels:

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

To configure basic VRRP for IPv6 support, configure VRRP group support on interfaces by including the `vrrp-inet6-group` statement:

```
vrrp-inet6-group group-id {  
    priority number;  
    virtual-inet6-address [ addresses ];  
    virtual-link-local-address ipv6-address;  
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address*]

Within a VRRP group, the master virtual router and the backup virtual router must be configured on two different routing platforms.

For each VRRP group, you must configure the following:

- Group identifier—Assign a value from 0 through 255.

- Address of one or more virtual routers that are members of the VRRP group—Normally, you configure only one virtual IP address per group. However, you can configure up to eight addresses. Do not include a prefix length in a virtual IP address. The following considerations apply to configuring a virtual IP address:
 - The virtual router IP address must be the same for all routing platforms in the VRRP group.
 - If you configure a virtual IP address to be the same as the physical interface's address, the interface becomes the master virtual router for the group. In this case, you must configure the priority to be 255, and you must configure preemption by including the **preempt** statement.
 - If the virtual IP address you choose is not the same as the physical interface's address, you must ensure that the virtual IP address does not appear anywhere else in the routing platform's configuration. Verify that you do not use this address for other interfaces, for the IP address of a tunnel, or for the IP address of static ARP entries.
 - You cannot configure a virtual IP address to be the same as the interface's address for an aggregated Ethernet interface. This configuration is not supported.
 - For VRRP for IPv6, the EUI-64 option cannot be used. In addition, the Duplicate Address Detection (DAD) process will not run for virtual IPv6 addresses.
 - You cannot configure the same virtual IP address on interfaces that belong to the same logical system and routing instance combination. However, you can configure the same virtual IP address on interfaces that belong to different logical system and routing instance combinations.
- Virtual link-local address—(VRRP for IPv6 only) You must explicitly define a virtual link-local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link-local address must be on the same subnet as the physical interface address.

- Priority for this routing platform to become the master virtual router—Configure the value used to elect the master virtual router in the VRRP group. It can be a number from 1 through 255. The default value for backup routers is 100. A larger value indicates a higher priority. The routing platform with the highest priority within the group becomes the master router.



NOTE: If there are two or more backup routers with the same priority, the router that has the highest primary address becomes the master.



NOTE: Mixed tagging (configuring two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing) is supported only for interfaces on Gigabit Ethernet IQ2 and IQ PICs. If you include the `flexible-vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level for a VRRP-enabled interface on a PIC that does not support mixed tagging, VRRP on that interface is disabled. In the output of the `show vrrp summary` command, the interface status is listed as Down.



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 level. (For more information, see the Junos® OS Network Interfaces.) 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 2378. 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 55](#)
- [Junos OS Support for VRRPv3 on page 56](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring the Startup Period for VRRP Operations on page 179](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 182](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 184](#)
- [Example: Configuring VRRP on page 197](#)

Configuring VRRP Authentication (IPv4 Only)

VRRP (IPv4 only) protocol exchanges can be authenticated to guarantee that only trusted routing platforms participate in routing in an autonomous system (AS). By default, VRRP authentication is disabled. You can configure one of the following authentication methods. Each VRRP group must use the same method.

- Simple authentication—Uses a text password included in the transmitted packet. The receiving routing platform uses an authentication key (password) to verify the packet.
- Message Digest 5 (MD5) algorithm—Creates the authentication data field in the IP authentication header. This header is used to encapsulate the VRRP PDU. The receiving routing platform uses an authentication key (password) to verify the authenticity of the IP authentication header and VRRP PDU.

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 routing platforms in the VRRP group.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]
- [edit logical-systems *logical-system-name* 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 (" "). The key must be the same for all routing platforms in the VRRP group.

You can include this statement at the following hierarchy levels:

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



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups. Therefore, if authentication is required, you need to configure alternative non-VRRP authentication mechanisms.

Related Documentation

- [Understanding VRRP on page 55](#)
- [Junos OS Support for VRRPv3 on page 56](#)
- [VRRP Configuration Hierarchy](#)

- [Configuring Basic VRRP Support on page 180](#)
- [Example: Configuring VRRP on page 197](#)

Configuring the Advertisement Interval for the VRRP Master Router

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

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.

For VRRP for IPv6, you must configure IPv6 router advertisements for the interface on which VRRP is configured to send IPv6 router advertisements for the VRRP group. To do so, include the **interface *interface-name*** statement at the **[edit protocols router-advertisement]** hierarchy level. (For information about this statement and guidelines, see the Junos OS Routing Protocols Configuration Guide.) When an interface receives an IPv6 router solicitation message, it sends an IPv6 router advertisement to all VRRP groups configured on it. In the case of logical systems, IPv6 router advertisements are not sent to VRRP groups.



NOTE: The master VRRP for an IPv6 router must respond to a router solicitation message with the virtual IP address of the router. However, when the **interface *interface-name*** statement is included at the **[edit protocols router-advertisement]** hierarchy level, the backup VRRP for an IPv6 router might send a response before the VRRP master responds, so that the default route of the client is not set to the master VRRP router's virtual IP address. To avoid this situation, include the **virtual-router-only** statement at the **[edit protocols router-advertisement interface *interface-name*]** hierarchy level. When this statement is included, router advertisements are sent only for VRRP IPv6 groups configured on the interface (if the groups are in the master state). You must include this statement on both the master and backup VRRP for IPv6 routers.

This topic contains the following sections:

- [Modifying the Advertisement Interval in Seconds on page 184](#)
- [Modifying the Advertisement Interval in Milliseconds on page 185](#)

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 levels:

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



NOTE: When VRRPv3 is enabled, the `advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

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 10 through 40,950 milliseconds.

You can include this statement at the following hierarchy levels:

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



NOTE: In the VRRP PDU, Junos OS sets the advertisement interval to 0. When you configure VRRP with other vendors' routers, the `fast-interval` statement works correctly only when the other routers also have an advertisement interval set to 0 in the VRRP PDUs. Otherwise, Junos OS interprets other routers' settings as advertisement timer errors.

To modify the time, in milliseconds, between the sending of VRRP for IPv6 advertisement packets, include the `inet6-advertise-interval` statement:

`inet6-advertise-interval ms;`

The range of values is from 100 through 40,950 milliseconds (ms).

You can include this statement at the following hierarchy levels:

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



NOTE: When VRRPv3 is enabled, the `inet6-advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

**Related
Documentation**

- [Understanding VRRP on page 55](#)
- [Junos OS Support for VRRPv3 on page 56](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring Basic VRRP Support on page 180](#)
- [Configuring a Backup Router to Preempt the Master Router on page 186](#)
- [Modifying the Preemption Hold-Time Value on page 187](#)
- [Configuring Asymmetric Hold Time for VRRP Routers on page 187](#)
- [Configuring the Silent Period on page 193](#)
- [Example: Configuring VRRP on page 197](#)

Configuring a Backup Router to Preempt the Master Router

By default, a higher-priority backup router preempts a lower-priority master router. To explicitly enable the master router to be preempted, include the `preempt` statement:

`preempt;`

You can include this statement at the following hierarchy levels:

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

To prohibit a higher-priority backup router from preempting a lower-priority master router, include the `no-preempt` statement:

`no-preempt;`

**Related
Documentation**

- [Understanding VRRP on page 55](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 184](#)
- [Modifying the Preemption Hold-Time Value on page 187](#)
- [Configuring Asymmetric Hold Time for VRRP Routers on page 187](#)
- [Example: Configuring VRRP on page 197](#)

Modifying the Preemption Hold-Time Value

The hold time is the maximum number of seconds that can elapse before a higher-priority backup router preempts the master router. You might want to configure a hold time so that all Junos OS components converge before preemption.

By default, the hold-time value is 0 seconds. A value of 0 means that preemption can occur immediately after the backup router comes online. Note that the hold time is counted from the time the backup router comes online. The hold time is only valid when the VRRP router is just coming online.

To modify the preemption hold-time value, include the **hold-time** statement:

```
hold-time seconds;
```

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy levels:

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

Related Documentation

- VRRP Configuration Hierarchy
- [Configuring the Advertisement Interval for the VRRP Master Router on page 184](#)
- [Configuring a Backup Router to Preempt the Master Router on page 186](#)
- [Configuring Asymmetric Hold Time for VRRP Routers on page 187](#)
- [Example: Configuring VRRP on page 197](#)

Configuring Asymmetric Hold Time for VRRP Routers

In Junos OS Release 9.5 and later, the **asymmetric-hold-time** statement at the [edit protocols vrrp] hierarchy level enables you to configure a VRRP master router to switch over to the backup router immediately—that is, without waiting for the priority hold time to expire—when a tracked interface or route goes down or when the bandwidth of a tracked interface decreases. Such events can cause an immediate reduction in the priority based on the configured priority cost for the event, and trigger a mastership election.

However, when the tracked route or interface comes up again, or when the bandwidth for a tracked interface increases, the backup (original master) router waits for the hold time to expire before it updates the priority and initiates the switchover if the priority is higher than the priority for the VRRP master (original backup) router.

If the **asymmetric-hold-time** statement is not configured, the VRRP master waits for the hold time to expire before it initiates a switchover when a tracked route goes down.

**Example: Configuring
Asymmetric Hold Time**

```
[edit]
user@host# set protocols vrrp asymmetric-hold-time
[edit]
user@host# show protocols vrrp
asymmetric-hold-time;
```

**Related
Documentation**

- VRRP Configuration Hierarchy
- [Configuring the Advertisement Interval for the VRRP Master Router on page 184](#)
- [Configuring a Backup Router to Preempt the Master Router on page 186](#)
- [Modifying the Preemption Hold-Time Value on page 187](#)
- [Example: Configuring VRRP on page 197](#)

Configuring an Interface to Accept Packets Destined for the Virtual IP Address

In VRRP implementations where the router acting as the master router is not the IP address owner—the IP address owner is the router that has the interface whose actual IP address is used as the virtual router's IP address (virtual IP address)—the master router accepts only the ARP packets from the packets that are sent to the virtual IP address. Junos OS enables you to override this limitation with the help of the **accept-data** configuration. When the **accept-data** statement is included in the configuration, the master router accepts all packets sent to the virtual IP address even when the master router is not the IP address owner.



NOTE: If the master router is the IP address owner or has its priority set to 255, the master router, by default, accepts all packets addressed to the virtual IP address. In such cases, the **accept-data** configuration is not required.

To configure an interface to accept all packets sent to the virtual IP address, include the **accept-data** statement:

```
accept-data;
```

You can include this statement at the following hierarchy levels:

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

To prevent a master router that is the IP address owner or has its priority set to 255 from accepting packets other than the ARP packets addressed to the virtual IP address, include the **no-accept-data** statement:

```
no-accept-data;
```

**NOTE:**

- 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, *Virtual Router Redundancy Protocol (VRRP)*).

Related Documentation

- [Understanding VRRP on page 55](#)
- [VRRP Configuration Hierarchy](#)
- [Example: Configuring VRRP on page 197](#)

Configuring a Logical Interface to Be Tracked

VRRP can track whether a logical interface is up, down, or not present, and can also dynamically change the priority of the VRRP group based on the state of the tracked logical interface, triggering a new master router election. VRRP can also track the operational speed of a logical interface and dynamically update the priority of the VRRP group when the speed crosses a configured threshold.

When interface tracking is enabled, you cannot configure a priority of 255 (a priority of 255 designates the master router). For each VRRP group, you can track up to 10 logical interfaces.

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 levels:

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

The interface specified is the interface to be tracked for the VRRP group. The priority hold time is the minimum length of time that must elapse between dynamic priority changes. A tracking event, such as an interface state change (up or down) or a change in bandwidth, triggers one of the following responses:

- The first tracking event initiates the priority hold timer, and also initializes the pending priority based on the current priority and the priority cost. However, the current priority remains unchanged.
- A tracking event or a manual configuration change that occurs while the priority hold timer is on triggers a pending priority update. However, the current priority remains unchanged.

This ensures that Junos OS does not initiate mastership elections every time a tracked interface flaps.

When the priority hold time expires, the current priority inherits the value from the pending priority, and the pending priority ceases.



NOTE: If you have configured `asymmetric-hold-time`, VRRP does not wait for the priority hold time to expire before initiating mastership elections if a tracked interface fails (state changes from up to down), or if the available bandwidth for a tracked interface decreases. For more information about `asymmetric-hold-time`, see [“Configuring Asymmetric Hold Time for VRRP Routers” on page 187](#).

The bandwidth threshold specifies a threshold for the tracked interface. When the bandwidth of the tracked interface drops below the configured bandwidth threshold value, the VRRP group uses the bandwidth threshold priority cost. You can track up to five bandwidth threshold statements for each tracked interface.

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked logical interface goes down, forcing a new master router election. The value can be from 1 through 254. The sum of the costs for all tracked logical interfaces or routes must be less than or equal to the configured priority of the VRRP group.

If you are tracking more than one interface, the router applies the sum of the priority costs for the tracked interfaces (at most, only one priority cost for each tracked interface) to the VRRP group priority. However, the interface priority cost and bandwidth threshold priority cost values for each VRRP group are not cumulative. The router uses only one priority cost to a tracked interface as indicated in [Table 18 on page 190](#):

Table 18: Interface State and Priority Cost Usage

Tracked Interface State	Priority Cost Usage
Down	<code>priority-cost priority</code>
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 have configured no bandwidth thresholds. If you have not configured an interface priority cost value, and the interface is down, the interface uses the bandwidth threshold priority cost value of the lowest bandwidth threshold.

Related Documentation

- [Understanding VRRP on page 55](#)
- VRRP Configuration Hierarchy
- [Configuring a Route to Be Tracked on page 191](#)
- [Example: Configuring VRRP on page 197](#)

Configuring a Route to Be Tracked

VRRP can track whether a route is reachable (that is, the route exists in the routing table of the routing instance included in the configuration) and dynamically change the priority of the VRRP group based on the reachability of the tracked route, triggering a new master router election.

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

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

You can include these statements at the following hierarchy levels:

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

The route prefix specified is the route to be tracked for the VRRP group. The priority hold time is the minimum length of time that must elapse between dynamic priority changes. A route tracking event, such as adding a route to or removing a route from the routing table, might trigger one or more of the following:

- The first tracking event initiates the priority hold timer, and also initializes the pending priority based on the current priority and the priority cost. However, the current priority remains unchanged.
- A tracking event or a manual configuration change that occurs while the priority hold timer is on triggers a pending priority update. However, the current priority remains unchanged.

When the priority hold time expires, the current priority inherits the value from the pending priority, and the pending priority ceases.

This ensures that Junos OS does not initiate mastership elections every time a tracked route flaps.



NOTE: If you have configured `asymmetric-hold-time`, VRRP does not wait for the priority hold time to expire before initiating mastership elections if a tracked route is removed from the routing table. For more information about `asymmetric-hold-time`, see [“Configuring Asymmetric Hold Time for VRRP Routers” on page 187](#).

The routing instance is the routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, specify the instance name as **default**.



NOTE: Tracking a route that belongs to a routing instance from a different logical system is not supported.

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked route goes down, forcing a new master router election. The value can be from 1 through 254. The sum of the costs for all tracked logical interfaces or routes must be less than or equal to the configured priority of the VRRP group.

**Related
Documentation**

- [Understanding VRRP on page 55](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring a Logical Interface to Be Tracked on page 189](#)
- [Example: Configuring VRRP Route Tracking on page 200](#)

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*. 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 frequent VRRP advertisements, and processes incoming VRRP advertisements. The groups that are inheriting the state do not process any incoming VRRP advertisement because the state is always inherited from the active VRRP group. However, the groups that are inheriting the state do send out VRRP advertisements once every 2 to 3 minutes to facilitate MAC address learning on the switches placed between the VRRP routers.

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*]** hierarchy level.

```
[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, the inheriting groups and the active group must be on the same physical interface and logical system. However, the groups do not need to necessarily be on the same routing instance (as was in Junos OS releases earlier than 9.6), VLAN, or logical interface.

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 55](#)
- [VRRP Configuration Hierarchy](#)

Configuring the Silent Period

The silent period starts when the interface state is changed from down to up. During this period, the Master Down Event is ignored. Configure the silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets during the interface startup phase.

To configure the silent period interval that the Master Down Event timer ignores, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

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



NOTE: During the silent startup period, the `show vrrp detail` command output shows a value of 0 for Master priority and your IP address for Master router. These values indicate that the Master selection is not completed yet, and these values can be ignored.

When you have configured **startup-silent-period**, the Master Down Event is ignored until the **startup-silent-period** expires.

For example, configure a VRRP group, *vrrp-group1*, with an advertise interval of 1 second, startup silent period of 10 seconds, and an interface *interface1* with a priority less than 255.

When *interface1* transitions from down to up:

- The *vrrp-group1* group moves to the backup state, and starts the Master Down Event timer (3 seconds; three times the value of the advertise interval, which is 1 second in this case).
- If no VRRP PDU is received during the 3-second period, the **startup-silent-period** (10 seconds in this case) is checked, and if the startup silent period has not expired, the Master Down Event timer is restarted. This is repeated until the **startup-silent-period** expires. In this example, the Master Down Event timer runs four times (12 seconds) by the time the 10-second startup silent period expires.
- If no VRRP PDU is received by the end of the fourth 3-second cycle, *vrrp-group1* takes over mastership.

**Related
Documentation**

- [Understanding VRRP on page 55](#)
- VRRP Configuration Hierarchy
- [startup-silent-period on page 221](#)

Configuring Passive ARP Learning for Backup VRRP Routers

By default, the backup VRRP router drops ARP requests for the VRRP-IP to VRRP-MAC address translation. This means that the backup router does not learn the ARP (IP-to-MAC address) mappings for the hosts sending the requests. When it detects a failure of the master router and transitions to become the new master router, the backup router must learn all the entries that were present in the ARP cache of the master router. In environments with many directly attached hosts, such as metro Ethernet environments, the number of ARP entries to learn can be high. This can cause a significant transition delay, during which the traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup router to hold approximately the same contents as the ARP cache in the master router, thus preventing the problem of learning ARP entries in a burst. 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 routers. Doing so prevents the need to manually intervene when the master router becomes the backup router. While a router is operating as the master router, the passive learning configuration has no operational impact. The configuration takes effect only when the router is operating as a backup router.

For information about configuring gratuitous ARP and the ARP aging timer, see the Junos OS System Basics Configuration Guide.

- Related Documentation**
- [Understanding VRRP on page 55](#)
 - VRRP Configuration Hierarchy

Enabling the Distributed Periodic Packet Management Process for VRRP

Typically, VRRP advertisements are sent by the VRRP process (vrripd) on the master VRRP router at regular intervals to let other members of the group know that the VRRP master router is operational.

When the vrripd process is busy and does not send VRRP advertisements, the backup VRRP routers might assume that the master router is down and take over as the master router, causing unnecessary flaps. This takeover might occur even though the original master router is still active and available, and might resume sending advertisements after the traffic has decreased. To address this problem and to reduce the load on the vrripd process, Junos OS uses the periodic packet management process (ppmd) to send VRRP advertisements on behalf of the vrripd process. However, you can further delegate the job of sending VRRP advertisements to the distributed ppmmd process that resides on the Packet Forwarding Engine.

The ability to delegate the sending of VRRP advertisements to the distributed ppmmd process ensures that the VRRP advertisements are sent even when the ppmmd process—which is now responsible for sending VRRP advertisements—is busy. Such delegation prevents the possibility of false alarms when the ppmmd process is busy. The ability to delegate the sending of VRRP advertisements to distributed ppmmd also adds to scalability because the load is shared across multiple ppmmd instances and is not concentrated on any single unit.



NOTE: CPU-intensive VRRP advertisements, such as advertisements with MD5 authentication or those sent and received over logical interfaces, such as Aggregated Ethernet interfaces, continue to be processed by the VRRP process on the Routing Engine even when distributed ppmmd is enabled.

To configure the distributed ppmmd process to send VRRP advertisements, include the **delegate-processing** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  delegate-processing;
```

- Related Documentation**
- [Understanding VRRP on page 55](#)
 - VRRP Configuration Hierarchy

Configuring VRRP to Improve Convergence Time

You can enable faster convergence time for the configured Virtual Router Redundancy Protocol (VRRP), thereby reducing the traffic restoration time to less than 1 second. To improve the convergence time for VRRP, perform the following tasks.

Before you begin, configure VRRP. See [“Example: Configuring VRRP” on page 197](#).

1. Configure the distributed periodic packet management (PPM) process to send VRRP advertisements when the VRRP process is busy.

```
[edit]
user@host# set protocols vrrp delegate-processing
```

2. Disable the skew timer to reduce the time required to transition to the master state.

```
[edit]
user@host# set protocols vrrp skew-timer-disable
```



NOTE: When there is only one master router and one backup router in the network deployment, you can disable the skew timer, thereby reducing the time required to transition to the master state.

3. Configure the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state.

```
[edit]
user@host# set protocols vrrp global-advertisement-threshold advertisement-value
```

4. Configure VRRP groups on the various subnets of a VLAN to inherit the state and to configure one of the groups.

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

5. Verify the configuration in operational mode.

```
[edit]
user@host> show protocols vrrp
```

**NOTE:**

- The reduction in convergence time is not applicable when VRRP is configured over integrated routing and bridging (IRB) interfaces, aggregated Ethernet interfaces, and multichassis link aggregation group (MC-LAG) interfaces.
- Compared to other routers, the convergence time and the traffic restoration time are less for MX Series routers with MPCs.
- Reduction in convergence time is applicable for all types of configurations at the physical interface, but the convergence time might not be less than 1 second for all the configurations. The convergence time depends on the number of groups that are transitioning from the backup to the master state and the interval at which these groups are transitioning.

Related Documentation

- [Improving the Convergence Time for VRRP on page 59](#)
- [Configuring Inheritance for a VRRP Group on page 192](#)
- [delegate-processing on page 210](#)
- [global-advertisements-threshold on page 212](#)
- [skew-timer-disable on page 221](#)

Example: Configuring VRRP

Configure one master (Router A) and one backup (Router B) routing platform. The address configured in the **virtual-address** statements differs from the addresses configured in the **address** statements. When you configure multiple VRRP groups on an interface, you configure one to be the master virtual router for that group.

On Router A

```
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.1.20/24 {
        vrrp-group 27 {
          virtual-address 192.168.1.15;
          priority 254;
          authentication-type simple;
          authentication-key booJUM;
        }
      }
    }
  }
}
```

On Router B

```
[edit interfaces]
ge-4/2/0 {
  unit 0 {
    family inet {
```

Configuring One Router to Be the Master Virtual Router for the Group

```

address 192.168.1.24/24 {
  vrrp-group 27 {
    virtual-address 192.168.1.15;
    priority 200;
    authentication-type simple;
    authentication-key booJUM;
  }
}

```

```

[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.1.20/24 {
        vrrp-group 2 {
          virtual-address 192.168.1.20;
          priority 255;
          advertise-interval 3;
          preempt;
        }
        vrrp-group 10 {
          virtual-address 192.168.1.55;
          priority 201;
          advertise-interval 3;
        }
        vrrp-group 1 {
          virtual-address 192.168.1.54;
          priority 22;
          advertise-interval 4;
        }
      }
    }
  }
}

```

Configuring VRRP and MAC Source Address Filtering

The VRRP group number is the decimal equivalent of the last byte of the virtual MAC address.

```

[edit interfaces]
ge-5/2/0 {
  gigether-options {
    source-filtering;
    source-address-filter {
      00:00:5e:00:01:0a; # Virtual MAC address
    }
  }
  unit 0 {
    family inet {
      address 192.168.1.10/24 {
        vrrp-group 10; # VRRP group number
        virtual-address 192.168.1.10;
        priority 255;
        preempt;
      }
    }
  }
}

```



```

    }
  }
}

```

- Related Documentation**
- [Understanding VRRP on page 55](#)
 - [VRRP Configuration Hierarchy](#)
 - [VRRP for IPv6 Configuration Hierarchy](#)
 - [Example: Configuring VRRP for IPv6 on page 199](#)
 - [Example: Configuring VRRP Route Tracking on page 200](#)

Example: Configuring VRRP for IPv6

Configure VRRP properties for IPv6 in one master (Router A) and one backup (Router B).

On Router A

```

[edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet6 {
      address fe80::5:0:0:6/64;
      address fec0::5:0:0:6/64 {
        vrrp-inet6-group 3; # VRRP inet6 group number
        virtual-inet6-address fec0::5:0:0:7;
        virtual-link-local-address fe80::5:0:0:7;
        priority 200;
        preempt;
      }
    }
  }
}

[edit protocols]
router-advertisement {
  interface ge-1/0/0.0 {
    prefix fec0::/64;
    max-advertisement-interval 4;
  }
}

```

On Router B

```

[edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet6 {
      address fe80::5:0:0:8/64;
      address fec0::5:0:0:8/64 {
        vrrp-inet6-group 3; # VRRP inet6 group number
        virtual-inet6-address fec0::5:0:0:7;
        virtual-link-local-address fe80::5:0:0:7;
        priority 100;
        preempt;
      }
    }
  }
}

```

```
    }  
  }  
  
  [edit protocols]  
  router-advertisement {  
    interface ge-1/0/0.0 {  
      prefix fec0::/64;  
      max-advertisement-interval 4;  
    }  
  }  
}
```

Related Documentation

- [Understanding VRRP on page 55](#)
- VRRP Configuration Hierarchy
- VRRP for IPv6 Configuration Hierarchy
- [Example: Configuring VRRP on page 197](#)
- [Example: Configuring VRRP Route Tracking on page 200](#)

Example: Configuring VRRP Route Tracking

Configure Routers R1 and R2 to run VRRP. Configure static routes and a policy for exporting the static routes on Router R3. The VRRP routing instances on R2 track the routes that are advertised by R3.

On Router R1

```
[edit interfaces]  
ge-1/0/3 {  
  unit 0 {  
    vlan-id 1;  
    family inet {  
      address 200.100.50.2/24 {  
        vrrp-group 0 {  
          virtual-address 200.100.50.101;  
          priority 195;  
        }  
      }  
    }  
  }  
}
```

On Router R2

```
[edit interfaces]  
ge-1/0/1 {  
  unit 0 {  
    vlan-id 1;  
    family inet {  
      address 200.100.50.1/24 {  
        vrrp-group 0 {  
          virtual-address 200.100.50.101;  
          priority 200;  
          track {  
            route 59.0.58.153/32 routing-instance default priority-cost 5;  
            route 59.0.58.154/32 routing-instance default priority-cost 5;  
            route 59.0.58.155/32 routing-instance default priority-cost 5;
```

```

    }
  }
}
}
}

```

On Router R3

```

[edit]
policy-options {
  policy-statement static-policy {
    term term1 {
      then accept;
    }
  }
}
protocols {
  ospf {
    export static-policy;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
routing-options {
  static {
    route 59.0.0.153/32 next-hop 45.45.45.46;
    route 59.0.0.154/32 next-hop 45.45.45.46;
    route 59.0.0.155/32 next-hop 45.45.45.46;
  }
}

```

- Related Documentation**
- [Understanding VRRP on page 55](#)
 - VRRP Configuration Hierarchy
 - VRRP for IPv6 Configuration Hierarchy
 - [Configuring a Route to Be Tracked on page 191](#)
 - [Example: Configuring VRRP on page 197](#)
 - [Example: Configuring VRRP for IPv6 on page 199](#)

Tracing VRRP Operations

To trace VRRP operations, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level.

By default, VRRP logs the error, data carrier detect (DCD) configuration, and routing socket events in a file in the **/var/log** directory. By default, this file is named **/var/log/vrrpd**.

The default file size is 1 megabyte (MB), and three files are created before the first one gets overwritten.

To change the configuration of the logging file, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
traceoptions {
  file filename <files number> <match regular-expression> <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
flag flag;
```

You can specify the following VRRP tracing flags:

- **all**—Trace all VRRP operations.
- **database**—Trace all database changes.
- **general**—Trace all general events.
- **interfaces**—Trace all interface changes.
- **normal**—Trace all normal events.
- **packets**—Trace all packets sent and received.
- **state**—Trace all state transitions.
- **timer**—Trace all timer events.

**Related
Documentation**

- [Understanding VRRP on page 55](#)
- VRRP Configuration Hierarchy

Configuration Statements: VRRP

- [\[edit protocols vrrp\] Hierarchy Level on page 203](#)

[\[edit protocols vrrp\] Hierarchy Level](#)


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

```
protocols {
  vrrp {
    asymmetric-hold-time;
    delegate-processing;
    failover-delay milliseconds;
    global-advertisements-threshold advertisement-value;
    skew-timer-disable;
    startup-silent-period seconds;
    traceoptions {
      file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
    version-3;
  }
}
```


Related Documentation

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)
- [\[edit protocols\] Hierarchy Level](#)
- [Junos OS Hierarchy and RFC Reference](#)
- [Junos® OS Ethernet Interfaces](#)
- [Junos® OS Network Interfaces](#)

accept-data

Syntax	(accept-data no-accept-data);
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not 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>  <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>	
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 an Interface to Accept Packets Destined for the Virtual IP Address on page 188


advertise-interval

Syntax	<code>advertise-interval seconds;</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.
Description	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets. All routers in the VRRP group must use the same advertisement interval.
<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	<i>seconds</i> —Interval between advertisement packets. Range: 1 through 255 seconds Default: 1 second
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Advertisement Interval for the VRRP Master Router on page 184 • fast-interval on page 211 • inet6-advertise-interval on page 214 • version-3 on page 225


asymmetric-hold-time

Syntax	asymmetric-hold-time;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Enable the VRRP master router to switch over to the backup router immediately, without waiting for the priority hold time to expire, when a route goes down. However, when the route comes back online, the backup router that is acting as the master waits for the priority hold time to expire before switching the mastership back to the original master VRRP router.
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 Asymmetric Hold Time for VRRP Routers on page 187

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> vrp-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> vrp-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.
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>  <p>NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups.</p> </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 182 • Configuring VRRP Authentication (IPv4 Only) • authentication-type on page 208 • version-3 on page 225

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.
Description	<p>Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the authentication-key statement.</p> <p>All routers in the VRRP group must use the same authentication scheme and password.</p> <div><p>NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups.</p></div>
Options	<p><i>authentication</i>—Authentication scheme:</p> <ul style="list-style-type: none">• simple—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure.• md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme. <p>Default: none (no authentication is performed).</p>
Required Privilege Level	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 182• Configuring VRRP Authentication (IPv4 Only)• authentication-key on page 207• version-3 on page 225

bandwidth-threshold

Syntax	<code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
Options	<p><i>bits-per-second</i>—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.</p> <p>Range: 1 through 10000000000000 bits per second</p> <p><i>priority-cost priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down 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 on page 189 • Configuring a Logical Interface to Be Tracked


delegate-processing (VRRP)

Syntax	delegate-processing;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Configure the distributed ppmmd process to send VRRP advertisements.
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">• Enabling the Distributed Periodic Packet Management Process for VRRP on page 195

failover-delay

Syntax	failover-delay <i>milliseconds</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Configure the failover delay for VRRP and VRRP for IPv6 operations.
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">• Configuring VRRP and VRRP for IPv6

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>
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 on page 184 • Configuring the Advertisement Interval for the VRRP Master • advertise-interval on page 205 • advertise-interval on page 205 • inet6-advertise-interval on page 214 • version-3 on page 225

global-advertisements-threshold

Syntax	global-advertisements-threshold <i>advertisement-value</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Configure the number of fast advertisements that can be missed by a backup router before the master router is declared as down.



NOTE:


- The advertisement value configured using the **global-advertisements-threshold** statement is applicable to all the Virtual Router Redundancy Protocol (VRRP) groups in the system.
 - Setting the advertisement value of the **global-advertisements-threshold** configuration to 1 is not recommended for a scaled configuration with an aggressive advertisement interval. For example, if you have 1000 VRRP groups with an advertisement interval of 100 ms, then do not set the **global-advertisements-threshold** value to 1.
 - Changing the advertisement value of the **global-advertisements-threshold** configuration during runtime can result in unpredictable behavior by the VRRP state machine. For example, momentary ownership change from the master router to the backup router and vice versa. Therefore, avoid changing the advertisement value of the **global-advertisements-threshold** statement during runtime.
-

Options	<i>advertisement-value</i> —Number of VRRP advertisements missed before the master router is declared as down. Range: 1 through 15 Default: 3
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">• Improving the Convergence Time for VRRP on page 59• Configuring VRRP to Improve Convergence Time on page 196

hold-time (VRRP)

Syntax	<code>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> vrrp-group group-id preempt],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id preempt],</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 group-id preempt],</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 group-id preempt]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	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	<p>seconds—Hold-time period.</p> <p>Range: 0 through 3600 seconds</p> <p>Default: 0 seconds (VRRP preemption is not timed.)</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 Backup Router to Preempt the Master Router on page 186 • Configuring VRRP Preemption and Hold Time


inet6-advertise-interval

Syntax	<code>inet6-advertise-interval <i>milliseconds</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id], [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 group-id]
Release Information	Statement introduced in Junos OS Release 8.4R2.
Description	<p>Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
	<div> NOTE: When VRRPv3 is enabled, the <code>inet6-advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</div>
Options	<p><i>milliseconds</i>—Interval, in milliseconds, between advertisement packets.</p> <p>Range: 100 to 40,000 milliseconds (ms)</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">• Configuring the Advertisement Interval for the VRRP Master Router on page 184• advertise-interval on page 205• fast-interval on page 211• version-3 on page 225

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>
Description	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p><i>interface-name</i>—Interface to be tracked for this VRRP group.</p> <p>Range: 1 through 10 interfaces</p> <p>The remaining statements are 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 on page 189 • Configuring a Logical Interface to Be Tracked • Junos Services Interfaces Configuration Release 11.2

preempt (VRRP)

Syntax	(preempt no-preempt) { hold-time seconds; }
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 interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-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>], [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>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, 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. <div style="margin-top: 10px;">  <p>NOTE: By default, a higher-priority backup router can preempt a lower-priority master router.</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"> • Configuring a Backup Router to Preempt the Master Router on page 186 • Configuring VRRP Preemption and Hold Time


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>
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 on page 180 • Configuring Basic VRRP Support

priority-cost (VRRP)

Syntax	<code>priority-cost priority;</code>
Hierarchy Level	<code>[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	priority —The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group. Range: 1 through 254
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Logical Interface to Be Tracked on page 189• Configuring a Logical Interface to Be Tracked


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>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.</p>
	<div>  <p>NOTE: When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP master will pre-empt according to the configuration of the priority-hold time.</p> </div>
Options	<p>seconds—Minimum length of time that must elapse between dynamic priority changes.</p> <p>Range: 0through 3600 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked on page 189 • Configuring a Logical Interface to Be Tracked

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 on page 191• Configuring a Route to Be Tracked

skew-timer-disable

Syntax	skew-timer-disable;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Disable the skew timer, thereby reducing the time required to transition from the backup state to the master state.
	<div>  <p>NOTE: The <code>skew-timer-disable</code> statement is used when there is only one master router and one backup router in the network.</p> </div>
Default	By default, the skew timer is enabled for all the VRRP groups.
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"> • Improving the Convergence Time for VRRP on page 59 • Configuring VRRP to Improve Convergence Time on page 196

startup-silent-period

Syntax	startup-silent-period <i>seconds</i> ;
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.
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	<p><i>seconds</i>—Number of seconds for the startup period.</p> <p>Default: 4 seconds</p> <p>Range: 1 through 2000 seconds</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 the Startup Period for VRRP Operations on page 179 • Configuring the Startup Period for VRRP Operations

traceoptions (Protocols VRRP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <microsecond-stamp> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; }</pre>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory /var/log.</p>
Default	If you do not include this statement, no VRRP-specific tracing operations are performed.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, VRRP tracing output is placed in the file vrrpd.</p> <p>files <i>number</i>—(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.</p> <p>Range: 0 through 4,294,967,296 files</p> <p>Default: 3 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. These are the VRRP-specific tracing options:</p> <ul style="list-style-type: none">• 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 *regular-expression*—(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—(Optional) 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—(Optional) Allow users to read the log file.

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

Related Documentation	<ul style="list-style-type: none">• Tracing VRRP Operations on page 201
------------------------------	---

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>
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 on page 189• Configuring a Route to Be Tracked on page 191• Configuring a Logical Interface to Be Tracked• Configuring a Route to Be Tracked

version-3

Syntax	version-3;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Enable Virtual Router Redundancy Protocol version 3 (VRRPv3).



NOTE:

- Even though the version-3 statement can be configured only at the [edit protocols vrrp] hierarchy level, VRRPv3 is enabled on all the configured logical systems as well.
 - When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with the previous versions of VRRP.
-

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"> • Junos OS Support for VRRPv3 on page 56 • VRRP Configuration Hierarchy • VRRP for IPv6 Configuration Hierarchy


virtual-address

Syntax	<code>virtual-address [<i>addresses</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>],</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>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX 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">• Configuring Basic VRRP Support on page 180• Configuring Basic VRRP Support

virtual-inet6-address

Syntax	<code>virtual-inet6-address [<i>addresses</i>];</code>
Hierarchy Level	<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>],</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>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) 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">• Configuring Basic VRRP Support on page 180

virtual-link-local-address

Syntax	<code>virtual-link-local-address <i>ipv6-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>], [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>]
Release Information	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Configure a virtual link-local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link-local address for each VRRP for IPv6 group. The virtual link-local address must be in the same subnet as the physical interface address.
	<div>  <p>NOTE: You do <i>not</i> need to configure link-local addresses and virtual link-local addresses when configuring VRRP for IPv6. Junos OS automatically generates link-local addresses and virtual link-local addresses. However, if link local addresses and virtual link-local addresses are configured, Junos OS considers the configured addresses.</p> </div>
Options	<i>ipv6-address</i> —virtual link-local IPv6 address for VRRP for an IPv6 group. Range: 0 through 255
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 Basic VRRP Support on page 180 • Junos OS Support for VRRPv3 on page 56

vrrp-group

Syntax	<pre> vrrp-group <i>group-id</i> { (<i>accept-data</i> <i>no-accept-data</i>); <i>advertise-interval seconds</i>; <i>advertisements-threshold number</i>; <i>authentication-key key</i>; <i>authentication-type authentication</i>; <i>fast-interval milliseconds</i>; (<i>preempt</i> <i>no-preempt</i>) { <i>hold-time seconds</i>; } <i>priority number</i>; track { interface <i>interface-name</i> { <i>bandwidth-threshold bits-per-second</i> <i>priority-cost priority</i>; <i>priority-cost priority</i>; } <i>priority-hold-time seconds</i>; route <i>prefix/prefix-length</i> <i>routing-instance instance-name</i> <i>priority-cost priority</i>; } <i>virtual-address [addresses]</i>; <i>vrrp-inherit-from vrrp-group</i>; } </pre>
Hierarchy Level	<pre> [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>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group.
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.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support on page 180 • Example: Configuring VRRP on page 197 • Configuring Basic VRRP Support

- Example: Configuring VRRP for Load Sharing
- [vrrp-inet6-group on page 229](#)

vrrp-inet6-group

Syntax `vrrp-inet6-group group-id {
 (accept-data | no-accept-data);
 advertisements-threshold number;
fast-interval milliseconds;
inet6-advertise-interval seconds;
 (preempt | no-preempt) {
hold-time seconds;
 }
priority number;
track {
interface interface-name {
bandwidth-threshold bits-per-second priority-cost priority;
priority-cost priority;
 }
priority-hold-time seconds;
route prefix/prefix-length routing-instance instance-name priority-cost priority;
 }
virtual-inet6-address [addresses];
virtual-link-local-address ipv6-address;
vrrp-inherit-from vrrp-group;
}`

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

Release Information Statement introduced before Junos OS Release 7.4.

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

Options *group-id*—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Range: 0 through 255

The remaining statements are explained separately.

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

Related Documentation

- [Configuring Basic VRRP Support on page 180](#)
- VRRP for IPv6 Configuration Hierarchy

PART 3

Administration

- [Routine Monitoring on page 233](#)
- [Operational Commands on page 239](#)

Routine Monitoring

- [Resetting Local Statistics on page 233](#)
- [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 234](#)
- [Verifying Graceful Restart Operation on page 234](#)

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 Junos OS Operational Mode Commands.



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 in the Junos OS on page 3](#)
- [Configuring Graceful Routing Engine Switchover on page 65](#)

Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart

Junos OS provides a tracing option to log restart signaling-based helper mode events for OSPF graceful restart. To enable tracing for restart signaling-based helper mode events, include the **traceoptions flag restart-signaling** statement at the **[edit protocols ospf]** hierarchy level.

To enable tracing for restart signaling-based events:

1. Create a log file for saving the log.

```
[edit protocols ospf]
user@host# set traceoptions file ospf-log
```

where *ospf-log* is the name of the log file.

2. Enable tracing for restart signaling-based helper mode events.

```
[edit protocols ospf]
user@host# set traceoptions flag restart-signaling
```

3. Commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

The logs are saved to the *ospf-log* file in the */var/log* folder.

Viewing the Log File

To view the restart signaling-based events from the log file, type:

```
user@host> file show /var/log/ospf-log | match "restart signaling"
Jun 25 14:44:08.890216 OSPF Restart Signaling: Start helper mode for nbr ip
14.19.3.2 id 10.10.10.1
Jun 25 14:44:11.358636 OSPF restart signaling: Received DBD with R bit set from
nbr ip=14.19.3.2 id=10.10.10.1. Start oob-resync.
Jun 25 14:44:11.380198 OSPF restart signaling: Received DBD with LR bit on from
nbr ip=14.19.3.2 id=10.10.10.1. Save its oob-resync capability 1
Jun 25 14:44:11.467200 OSPF restart signaling: nbr fsm for nbr ip=14.19.3.2
id=10.10.10.1 moving to state Full. Reset oob-resync parameters.
```

Related Documentation

- Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart
- [Example: Managing Helper Modes for OSPF Graceful Restart on page 116](#)
- helper-disable (OSPF)

Verifying Graceful Restart Operation

This topic contains the following sections:

- [Graceful Restart Operational Mode Commands on page 235](#)
- [Verifying BGP Graceful Restart on page 235](#)

- [Verifying IS-IS and OSPF Graceful Restart on page 236](#)
- [Verifying CCC and TCC Graceful Restart on page 236](#)

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 Junos OS Operational Mode Commands.

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 86](#)).

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)                Legend for connection types
UN -- uninitialized                    if-sw: interface switching
NP -- not present                     rmt-if: remote interface switching
WE -- wrong encapsulation              lsp-sw: LSP switching
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting
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	r1sp	Up	
CE1-CE2-3	rmt-if	Restart	----- 0
fe-1/1/0.3	intf	Up	
PE1-PE2-3	t1sp	Up	
PE2-PE1-3	r1sp	Up	

- Related Documentation**
- [Graceful Restart Concepts on page 29](#)
 - Configuring Graceful Restart for QFabric Systems

CHAPTER 20

Operational Commands

show system switchover

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.
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 routers (or line-card chassis) in the routing matrix.

Options	all-chassis —(TX Matrix 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 routers configured in the routing matrix.

all-lcc—(TX Matrix 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 T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.

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 and TX Matrix Plus router only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for a specific T640 router (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for a specific T1600 router (or line-card chassis) connected to the TX Matrix Plus router. Replace ***number*** with 0.

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

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

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

If you issue the **show system switchover** command on the active Routing Engine in the master router of an MX Series Virtual Chassis, the router displays an error message that graceful Routing Engine switchover (GRES) is not enabled on this member.

Required Privilege Level view

List of Sample Output [show system switchover \(Backup Routing Engine\) on page 243](#)
[show system switchover all-lcc \(Routing Matrix\) on page 243](#)

Output Fields Table 19 on page 242 describes the output fields for the **show system switchover** command. Output fields are listed in the approximate order in which they appear.

Table 19: show system switchover Output Fields

Field Name	Field Description
Graceful switchover	Display graceful Routing Engine switchover status: <ul style="list-style-type: none"> • On—Indicates graceful-switchover is specified for the routing-options configuration command. • Off—Indicates graceful-switchover is not specified for the routing-options configuration command.
Configuration database	State of the configuration database: <ul style="list-style-type: none"> • Ready—Configuration database has synchronized. • Synchronizing—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds. • Synchronize failed—Configuration database synchronize process failed.
Kernel database	State of the kernel database: <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 /var/log/ksyncd for possible causes, or notify Juniper Networks customer support.
Peer state	Routing Engine peer state: <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
```


PART 4

Troubleshooting

- [Troubleshooting Unified ISSU on page 247](#)

CHAPTER 21

Troubleshooting Unified ISSU

- [Troubleshooting Unified ISSU Problems on page 247](#)

Troubleshooting Unified ISSU Problems

If the unified ISSU procedure stops progressing, complete the following steps:

1. Open a new session on the master Routing Engine and issue the **request system software abort in-service-upgrade** command.
2. Check the existing router session to verify that the upgrade has been aborted.

An “ISSU: aborted!” message is provided. Additional system messages provide you with information about where the upgrade stopped and recommendations for the next step to take.

For more information about the **request system software abort in-service-upgrade** command, see the Junos OS Operational Mode Commands.

Related Documentation

- [Unified ISSU Concepts on page 37](#)
- [Unified ISSU Process on the TX Matrix Router](#)
- [Unified ISSU System Requirements on page 42](#)
- [Best Practices on page 153](#)
- [Performing a Unified ISSU on page 157](#)
- [Verifying a Unified ISSU on page 169](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 170](#)

