



---

Junos<sup>®</sup> OS

## High Availability Feature Guide for Routing Devices



---

Modified: 2018-05-29

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

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

*Junos® OS High Availability Feature Guide for Routing Devices*  
Copyright © 2018 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 <https://www.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

	About the Documentation . . . . .	xi
	Documentation and Release Notes . . . . .	xi
	Supported Platforms . . . . .	xi
	Using the Examples in This Manual . . . . .	xi
	Merging a Full Example . . . . .	xii
	Merging a Snippet . . . . .	xii
	Documentation Conventions . . . . .	xiii
	Documentation Feedback . . . . .	xv
	Requesting Technical Support . . . . .	xv
	Self-Help Online Tools and Resources . . . . .	xv
	Opening a Case with JTAC . . . . .	xvi
<b>Chapter 1</b>	<b>Overview . . . . .</b>	<b>17</b>
	High Availability Overview . . . . .	17
	Understanding High Availability Features on Juniper Networks Routers . . . . .	17
	Routing Engine Redundancy . . . . .	17
	Graceful Routing Engine Switchover . . . . .	18
	Nonstop Bridging . . . . .	18
	Nonstop Active Routing . . . . .	18
	Graceful Restart . . . . .	19
	Nonstop Active Routing Versus Graceful Restart . . . . .	20
	Effects of a Routing Engine Switchover . . . . .	21
	VRRP . . . . .	21
	Unified ISSU . . . . .	21
	Interchassis Redundancy for MX Series Routers Using Virtual Chassis . . . . .	22
	High Availability-Related Features in Junos OS . . . . .	23
<b>Chapter 2</b>	<b>Configuring Switching Control Board Redundancy . . . . .</b>	<b>25</b>
	Understanding Switching Control Board Redundancy . . . . .	25
	Redundant CFEBs on the M10i Router . . . . .	25
	Redundant FEBs on the M120 Router . . . . .	26
	Redundant SSBs on the M20 Router . . . . .	28
	Redundant SFMs on the M40e and M160 Routers . . . . .	28
	Configuring CFEB Redundancy on the M10i Router . . . . .	29
	Configuring Switching Control Board Redundancy . . . . .	30
	Configuring CFEB Redundancy on the M10i Router . . . . .	30
	Configuring FEB Redundancy on the M120 Router . . . . .	30
	Example: Configuring FEB Redundancy on M120 Routers . . . . .	31
	Configuring SFM Redundancy on M40e and M160 Routers . . . . .	33

	Configuring SSB Redundancy on the M20 Router . . . . .	34
	Configuring SFM Redundancy on M40e and M160 Routers . . . . .	34
	Configuring SSB Redundancy on the M20 Router . . . . .	35
<b>Chapter 3</b>	<b>Configuring Bidirectional Forwarding Detection (BFD) . . . . .</b>	<b>37</b>
	Configuring BFD for Static Routes . . . . .	37
	Understanding BFD for Static Routes for Faster Network Failure Detection . . . . .	37
	Understanding Static Route State When BFD is in Admin Down State . . . . .	42
	Example: Configuring BFD for Static Routes for Faster Network Failure Detection . . . . .	42
	Configuring BFD for BGP . . . . .	48
	Understanding BFD for BGP . . . . .	48
	Example: Configuring BFD on Internal BGP Peer Sessions . . . . .	50
	Configuring BFD for OSPF . . . . .	59
	Understanding BFD for OSPF . . . . .	59
	Example: Configuring BFD for OSPF . . . . .	61
	Configuring BFD for IS-IS . . . . .	65
	Understanding BFD for IS-IS . . . . .	66
	Example: Configuring BFD for IS-IS . . . . .	68
	Configuring BFD for RIP . . . . .	74
	Understanding BFD for RIP . . . . .	75
	Example: Configuring BFD for RIP . . . . .	75
	Configuring BFD for PIM . . . . .	81
	Configuring Independent Micro BFD Sessions for LAG . . . . .	83
	Understanding Independent Micro BFD Sessions for LAG . . . . .	83
	Configuring Independent Micro BFD Sessions for LAG . . . . .	85
	Example: Configuring Independent Micro BFD Sessions for LAG . . . . .	91
	Understanding Distributed BFD . . . . .	100
<b>Part 1</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 4</b>	<b>Configuration Statements: Bidirectional Forwarding Detection . . . . .</b>	<b>107</b>
	authentication (LAG) . . . . .	108
	bfd-liveness-detection (LAG) . . . . .	109
	detection-time (LAG) . . . . .	111
	traceoptions (Protocols BFD) . . . . .	112
	transmit-interval (LAG) . . . . .	114
<b>Chapter 5</b>	<b>Configuration Statements: Graceful Routing Engine Switchover . . . . .</b>	<b>115</b>
	graceful-switchover . . . . .	115
<b>Chapter 6</b>	<b>Configuration Statements: Graceful Restart . . . . .</b>	<b>117</b>
	disable . . . . .	118
	graceful-restart (Enabling Globally) . . . . .	119
	graceful-restart (Multicast Snooping) . . . . .	120
	helper-disable (Multiple Protocols) . . . . .	121
	helper-disable (OSPF) . . . . .	122
	maximum-helper-recovery-time . . . . .	123
	maximum-helper-restart-time (RSVP) . . . . .	124

	maximum-neighbor-reconnect-time . . . . .	125
	maximum-neighbor-recovery-time . . . . .	126
	no-strict-lsa-checking . . . . .	127
	notify-duration . . . . .	128
	not-on-disk-underperform . . . . .	129
	reconnect-time . . . . .	130
	recovery-time . . . . .	131
	restart-duration . . . . .	132
	restart-time (BGP Graceful Restart) . . . . .	133
	stale-routes-time . . . . .	134
	traceoptions (Protocols) . . . . .	135
<b>Chapter 7</b>	<b>Configuration Statements: Nonstop Active Routing . . . . .</b>	<b>137</b>
	nonstop-routing . . . . .	138
	switchover-on-routing-crash . . . . .	139
	synchronize . . . . .	140
	traceoptions . . . . .	142
<b>Chapter 8</b>	<b>Configuration Statements: Nonstop Bridging . . . . .</b>	<b>145</b>
	nonstop-bridging . . . . .	145
<b>Chapter 9</b>	<b>Configuration Statements: Routing Engine and Switching Control Board Redundancy . . . . .</b>	<b>147</b>
	cfeb . . . . .	148
	description (Chassis Redundancy) . . . . .	148
	failover (Chassis) . . . . .	149
	failover (System Process) . . . . .	150
	feb (Creating a Redundancy Group) . . . . .	151
	feb (Assigning a FEB to a Redundancy Group) . . . . .	152
	keepalive-time . . . . .	153
	no-auto-failover . . . . .	154
	on-disk-failure (Chassis Redundancy Failover) . . . . .	154
	on-loss-of-keepalives . . . . .	155
	redundancy . . . . .	156
	redundancy-group . . . . .	157
	routing-engine (Chassis Redundancy) . . . . .	158
	sfm (Chassis Redundancy) . . . . .	159
	ssb . . . . .	160
<b>Chapter 10</b>	<b>Configuration Statements: Unified ISSU . . . . .</b>	<b>161</b>
	no-issu-timer-negotiation . . . . .	161
	traceoptions (Protocols BFD) . . . . .	162
<b>Chapter 11</b>	<b>Configuration Statements: VRRP . . . . .</b>	<b>165</b>
	accept-data . . . . .	167
	advertise-interval . . . . .	168
	asymmetric-hold-time . . . . .	169
	authentication-key . . . . .	170
	authentication-type . . . . .	171
	bandwidth-threshold . . . . .	172
	delegate-processing (VRRP) . . . . .	173

	fast-interval . . . . .	174
	global-advertisements-threshold . . . . .	175
	hold-time (VRRP) . . . . .	176
	inherit-advertisement-interval . . . . .	177
	inet6-advertise-interval . . . . .	178
	interface . . . . .	179
	preempt (VRRP) . . . . .	180
	priority (Protocols VRRP) . . . . .	181
	priority-cost (VRRP) . . . . .	182
	priority-hold-time . . . . .	183
	route (Interfaces) . . . . .	184
	skew-timer-disable . . . . .	185
	startup-silent-period . . . . .	186
	traceoptions (Protocols VRRP) . . . . .	187
	track (VRRP) . . . . .	189
	version-3 . . . . .	190
	virtual-address . . . . .	191
	virtual-inet6-address . . . . .	192
	virtual-link-local-address . . . . .	193
	vrrp-group . . . . .	194
	vrrp-inet6-group . . . . .	196
	vrrp-inherit-from . . . . .	197
<b>Chapter 12</b>	<b>Operational Commands . . . . .</b>	<b>199</b>
	clear vrrp . . . . .	200
	request chassis ssb master switch . . . . .	201
	request system software in-service-upgrade . . . . .	203
	request system software in-service-upgrade (MX Series 3D Universal Edge Routers and EX9200 Switches) . . . . .	216
	request system software validate in-service-upgrade . . . . .	233
	show chassis ssb . . . . .	236
	show nonstop-routing . . . . .	238
	show pfe ssb . . . . .	241
	show system switchover . . . . .	247
	show task replication . . . . .	254
	show vrrp . . . . .	256
	show vrrp track . . . . .	267

# List of Figures

<b>Chapter 3</b>	<b>Configuring Bidirectional Forwarding Detection (BFD) . . . . .</b>	<b>37</b>
	Figure 1: Customer Routes Connected to a Service Provider . . . . .	43
	Figure 2: Typical Network with IBGP Sessions . . . . .	52
	Figure 3: Configuring BFD for IS-IS . . . . .	69
	Figure 4: RIP BFD Network Topology . . . . .	77
	Figure 5: Configuring an Independent Micro BFD Session for LAG . . . . .	91





# List of Tables

	<b>About the Documentation . . . . .</b>	<b>xi</b>
	Table 1: Notice Icons . . . . .	xiii
	Table 2: Text and Syntax Conventions . . . . .	xiv
<b>Chapter 3</b>	<b>Configuring Bidirectional Forwarding Detection (BFD) . . . . .</b>	<b>37</b>
	Table 3: Configuring BFD for IS-IS . . . . .	66
	Table 4: BFD Modes Supported on SRX Series Devices . . . . .	101
<b>Part 1</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 12</b>	<b>Operational Commands . . . . .</b>	<b>199</b>
	Table 5: show chassis ssb Output Fields . . . . .	236
	Table 6: show nonstop-routing Output Fields . . . . .	238
	Table 7: show pfe ssb Output Fields . . . . .	241
	Table 8: show system switchover Output Fields . . . . .	249
	Table 9: show task replication Output Fields . . . . .	254
	Table 10: show vrrp Output Fields . . . . .	257
	Table 11: show vrrp track Output Fields . . . . .	268



# About the Documentation

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

## Documentation and Release Notes

---

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

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 <https://www.juniper.net/books>.

## Supported Platforms

---

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

- M Series
- T Series
- MX Series
- TX Matrix
- PTX Series

## Using the Examples in This Manual

---

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming

configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

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

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

## Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

## Merging a Snippet

To merge a snippet, follow these steps:

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

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

```
commit {
  file ex-script-snippet.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 [CLI Explorer](#).

## Documentation Conventions

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

Table 1: Notice Icons







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

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

Table 2: Text and Syntax Conventions

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

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
GUI Conventions		
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

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

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

## Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

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

## 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 <https://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 <https://www.juniper.net/support/requesting-support.html>.



## CHAPTER 1

# Overview

- [High Availability Overview on page 17](#)

## High Availability Overview

---

- [Understanding High Availability Features on Juniper Networks Routers on page 17](#)
- [High Availability-Related Features in Junos OS on page 23](#)

## Understanding High Availability Features on Juniper Networks Routers

For Juniper Networks routing platforms running the Junos operating system (Junos OS), *high availability* refers to the hardware and software components that provide redundancy and reliability for packet-based communications. This topic provides brief overviews of the following high availability features:

- [Routing Engine Redundancy on page 17](#)
- [Graceful Routing Engine Switchover on page 18](#)
- [Nonstop Bridging on page 18](#)
- [Nonstop Active Routing on page 18](#)
- [Graceful Restart on page 19](#)
- [Nonstop Active Routing Versus Graceful Restart on page 20](#)
- [Effects of a Routing Engine Switchover on page 21](#)
- [VRRP on page 21](#)
- [Unified ISSU on page 21](#)
- [Interchassis Redundancy for MX Series Routers Using Virtual Chassis on page 22](#)

## Routing Engine Redundancy

---

Redundant Routing Engines are two Routing Engines that are installed in the same routing platform. One functions as the master, while the other stands by as a backup should the master Routing Engine fail. On routing platforms with dual Routing Engines, network reconvergence takes place more quickly than on routing platforms with a single Routing Engine.

## Graceful Routing Engine Switchover

---

Graceful Routing Engine switchover (GRES) 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.



**NOTE:** To preserve routing during a switchover, graceful Routing Engine switchover must be combined with either graceful restart protocol extensions or nonstop active routing. For more information, see *Understanding Graceful Routing Engine Switchover* and *Nonstop Active Routing Concepts*.



**NOTE:** In T Series routers, TX Matrix routers, and TX Matrix Plus routers, the control plane is preserved in case of GRES with NSR, and 75% of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES.

## Nonstop Bridging

---

Nonstop bridging enables an MX Series 3D Universal Edge Router with redundant Routing Engines to switch from a primary Routing Engine to a backup Routing Engine without losing Layer 2 Control Protocol (L2CP) information. Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover to preserve interface and kernel information. However, nonstop bridging also saves 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.

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)

For more information, see *Nonstop Bridging Concepts*.

## Nonstop Active Routing

---

Nonstop active routing (NSR) enables a routing platform with redundant Routing Engines to switch from a primary Routing Engine to a backup Routing Engine without alerting peer nodes that a change has occurred. Nonstop active routing uses the same infrastructure as graceful Routing Engine switchover to preserve interface and kernel

information. However, nonstop active routing also preserves routing information and protocol sessions by running the routing protocol process (rpd) on both Routing Engines. In addition, nonstop active routing preserves TCP connections maintained in the kernel.



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

For a list of protocols and features supported by nonstop active routing, see *Nonstop Active Routing Protocol and Feature Support*.

For more information about nonstop active routing, see *Nonstop Active Routing Concepts*.

### Graceful Restart

With routing protocols, any service interruption requires an affected router to 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. To alleviate this situation, graceful restart provides extensions to routing protocols. These protocol extensions define two roles for a router—*restarting* and *helper*. The extensions signal neighboring routers about a router undergoing a restart and prevent the neighbors from propagating the change in state to the network during a graceful restart wait interval. 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.

When a router is running graceful restart and the router stops sending and replying to protocol liveness messages (hellos), the adjacencies assume a graceful restart and begin running a timer to monitor the restarting router. During this interval, helper routers do not process an adjacency change for the router that they assume is restarting, but continue active routing with the rest of the network. The helper routers assume that the router can continue stateful forwarding based on the last preserved routing state during the restart.

If the router was actually restarting and is back up before the graceful timer period expires in all of the helper routers, the helper routers provide the router with the routing table, topology table, or label table (depending on the protocol), exit the graceful period, and return to normal network routing.

If the router does not complete its negotiation with helper routers before the graceful timer period expires in all of the helper routers, the helper routers process the router's change in state and send routing updates, so that convergence occurs across the network. If a helper router detects a link failure from the router, the topology change causes the helper router to exit the graceful wait period and to send routing updates, so that network convergence occurs.

To enable a router to undergo a graceful restart, you must include the **graceful-restart** statement at the global **[edit routing-options]** or **[edit routing-instances instance-name routing-options]** hierarchy level. You can, optionally, modify the global settings at the

individual protocol level. When a routing session is started, a router that is configured with graceful restart must negotiate with its neighbors to support it when it undergoes a graceful restart. A neighboring router will accept the negotiation and support helper mode without requiring graceful restart to be configured on the neighboring router.



**NOTE:** A Routing Engine switchover event on a helper router that is in graceful wait state causes the router to drop the wait state and to propagate the adjacency's state change to the network.

Graceful restart is supported for the following protocols and applications:

- BGP
- ES-IS
- IS-IS
- OSPF/OSPFv3
- PIM sparse mode
- RIP/RIPng
- MPLS-related protocols, including:
  - Label Distribution Protocol (LDP)
  - Resource Reservation Protocol (RSVP)
  - Circuit cross-connect (CCC)
  - Translational cross-connect (TCC)
- Layer 2 and Layer 3 virtual private networks (VPNs)

For more information, see *Graceful Restart Concepts*.

### **Nonstop Active Routing Versus Graceful Restart**

---

Nonstop active routing and graceful restart are two different methods of maintaining high availability. Graceful restart requires a router restart. A router undergoing a graceful restart relies on its neighbors (or helpers) to restore its routing protocol information. The restart is the mechanism by which helpers are signaled to exit the wait interval and start providing routing information to the restarting router. For more information, see *Graceful Restart Concepts*.

In contrast, nonstop active routing does not involve a router restart. Both the master and backup Routing Engines are running the routing protocol process (rpd) and exchanging updates with neighbors. When one Routing Engine fails, the router simply switches to the active Routing Engine to exchange routing information with neighbors. Because of these feature differences, nonstop routing and graceful restart are mutually exclusive. Nonstop active routing cannot be enabled when the router is configured as a graceful restarting router. If you include the **graceful-restart** statement at any hierarchy level and the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and try to

commit the configuration, the commit request fails. For more information, see *Nonstop Active Routing Concepts*.

### Effects of a Routing Engine Switchover

---

*Effects of a Routing Engine Switchover* 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.

### VRRP

---

The Virtual Router Redundancy Protocol (VRRP) enables hosts on a LAN to make use of redundant routing platforms (master and backup pairs) on the LAN, requiring only the static configuration of a single default route on the hosts.

The VRRP routing platform pairs 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.

VRRP has advantages in ease of administration and network throughput and reliability:

- It provides a virtual default routing platform.
- It enables traffic on the LAN to be routed without a single point of failure.
- A virtual backup router can take over a failed default router:
  - Within a few seconds.
  - With a minimum of VRRP traffic.
  - Without any interaction with the hosts.

Devices 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 (default 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.

As of Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the **nonstop-routing** statement at the **[edit routing-options]** or **[edit logical-system logical-system-name routing-options]** hierarchy level.

For more information, see *Understanding VRRP*.

### Unified ISSU

---

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 by dual Routing Engine platforms. In

addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

With a unified ISSU, you can eliminate network downtime, reduce operating costs, and deliver higher services levels. For more information, see *Getting Started with Unified In-Service Software Upgrade*.

### **Interchassis Redundancy for MX Series Routers Using Virtual Chassis**

---

*Interchassis redundancy* is a high availability feature that can span equipment located across multiple geographies to prevent network outages and protect routers against access link failures, uplink failures, and wholesale chassis failures without visibly disrupting the attached subscribers or increasing the network management burden for service providers. As more high-priority voice and video traffic is carried on the network, interchassis redundancy has become a requirement for providing stateful redundancy on broadband subscriber management equipment such as broadband services routers, broadband network gateways, and broadband remote access servers. Interchassis redundancy support enables service providers to fulfill strict service-level agreements (SLAs) and avoid unplanned network outages to better meet the needs of their customers.

To provide a stateful interchassis redundancy solution for MX Series 3D Universal Edge Routers, you can configure a Virtual Chassis. A *Virtual Chassis* configuration interconnects two MX Series routers into a logical system that you can manage as a single network element. The member routers in a Virtual Chassis are designated as the *master router* (also known as the *protocol master*) and the *backup router* (also known as the *protocol backup*). The member routers are interconnected by means of dedicated *Virtual Chassis ports* that you configure on Trio Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces.

An MX Series Virtual Chassis is managed by the *Virtual Chassis Control Protocol (VCCP)*, which is a dedicated control protocol based on IS-IS. VCCP runs on the Virtual Chassis port interfaces and is responsible for building the Virtual Chassis topology, electing the Virtual Chassis master router, and establishing the interchassis routing table to route traffic within the Virtual Chassis.

Starting with Junos OS Release 11.2, Virtual Chassis configurations are supported on MX240, MX480, and MX960 3D Universal Edge Routers with Trio MPC/MIC interfaces and dual Routing Engines. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled on both member routers in the Virtual Chassis.

**See Also**    • [High Availability-Related Features in Junos OS on page 23](#)

## High Availability-Related Features in Junos OS

Related redundancy and reliability features include:

- Redundant power supplies, host modules, host subsystems, and forwarding boards. For more information, see the *Junos OS Administration Library* and the *Junos OS Hardware Network Operations Guide*.
- Additional link-layer redundancy, including Automatic Protection Switching (APS) for SONET interfaces, Multiplex Section Protection (MSP) for SDH interfaces, and DLSw redundancy for Ethernet interfaces. For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.
- Bidirectional Forwarding Detection (BFD) works with other routing protocols to detect failures rapidly. For more information, see the *Junos OS Routing Protocols Library*.
- Redirection of Multiprotocol Label Switching (MPLS) label-switched path (LSP) traffic—Mechanisms such as link protection, node-link protection, and fast reroute recognize link and node failures, allowing MPLS LSPs to select a bypass LSP to circumvent failed links or devices. For more information, see the *MPLS Applications Feature Guide*.

**See Also** • [Understanding High Availability Features on Juniper Networks Routers on page 17](#)

**Related Documentation** •





## CHAPTER 2

# Configuring Switching Control Board Redundancy

- [Understanding Switching Control Board Redundancy on page 25](#)
- [Configuring CFEB Redundancy on the M10i Router on page 29](#)
- [Configuring Switching Control Board Redundancy on page 30](#)
- [Configuring SFM Redundancy on M40e and M160 Routers on page 34](#)
- [Configuring SSB Redundancy on the M20 Router on page 35](#)

## Understanding Switching Control Board Redundancy

---

This section describes the following redundant switching control boards:



**NOTE:** A failover from a master switching control board to a backup switching control board occurs automatically when the master experiences a hardware failure or when you have configured the software to support a change in mastership based on specific conditions. You can also manually switch mastership by issuing specific `request chassis` commands. In this section, the term *failover* refers to an automatic event, whereas *switchover* refers to either an automatic or a manual event.

- [Redundant CFEBs on the M10i Router on page 25](#)
- [Redundant FEBs on the M120 Router on page 26](#)
- [Redundant SSBs on the M20 Router on page 28](#)
- [Redundant SFMs on the M40e and M160 Routers on page 28](#)

## Redundant CFEBs on the M10i Router

On the M10i router, the CFEB performs the following functions:

- **Route lookups**—Performs route lookups using the forwarding table stored in synchronous SRAM (SSRAM).
- **Management of shared memory**—Uniformly allocates incoming data packets throughout the router's shared memory.

- Transfer of outgoing data packets—Passes data packets to the destination Fixed Interface Card (FIC) or Physical Interface Card (PIC) when the data is ready to be transmitted.
- Transfer of exception and control packets—Passes exception packets to the microprocessor on the CFEB, which processes almost all of them. The remainder are sent to the Routing Engine for further processing. Any errors originating in the Packet Forwarding Engine and detected by the CFEB are sent to the Routing Engine using system log messages.

The M10i router has two CFEBs, one that is configured to act as the master and the other that serves as a backup in case the master fails. You can initiate a manual switchover by issuing the **request chassis cfeb master switch** command. For more information, see the *Junos OS Administration Library*.

## Redundant FEBs on the M120 Router

The M120 router supports up to six Forwarding Engine Boards (FEBs). Flexible PIC Concentrator (FPCs), which host PICs, are separate from the FEBs, which handle packet forwarding. FPCs are located on the front of the chassis and provide power and management to PICs through the midplane. FEBs are located on the back of the chassis and receive signals from the midplane, which the FEBs process for packet forwarding. The midplane allows any FEB to carry traffic for any FPC.

To configure the mapping of FPCs to FEBs, use the **fpc-feb-connectivity** statement as described in the *Junos OS Administration Library*. You cannot specify a connection between an FPC and a FEB configured as a backup. If an FPC is not specified to connect to a FEB, the FPC is assigned automatically to the FEB with the same slot number. For example, the FPC in slot 1 is assigned to the FEB in slot 1.

You can configure one FEB as a backup for one or more FEBs by configuring a FEB redundancy group. When a FEB fails, the backup FEB can quickly take over packet forwarding. A redundancy group must contain exactly one backup FEB and can optionally contain one primary FEB and multiple other FEBs. A FEB can belong to only one group. A group can provide backup on a one-to-one basis (primary-to-backup), a many-to-one basis (two or more other-FEBs-to-backup), or a combination of both (one primary-to-backup and one or more other-FEBs-to-backup).

When you configure a primary FEB in a redundancy group, the backup FEB mirrors the exact forwarding state of the primary FEB. If switchover occurs from a primary FEB, the backup FEB does not reboot. A manual switchover from the primary FEB to the backup FEB results in less than 1 second of traffic loss. Failover from the primary FEB to the backup FEB results in less than 10 seconds of traffic loss.

If a failover occurs from the other FEB and a primary FEB is specified for the group, the backup FEB reboots so that the forwarding state from the other FEB can be downloaded to the backup FEB and forwarding can continue. Automatic failover from a FEB that is not specified as a primary FEB results in higher packet loss. The duration of packet loss depends on the number of interfaces and on the size of the routing table, but it can be minutes.

If a failover from a FEB occurs when no primary FEB is specified in the redundancy group, the backup FEB does not reboot and the interfaces on the FPC connected to the previously active FEB remain online. The backup FEB must obtain the entire forwarding state from the Routing Engine after a switchover, and this update may take a few minutes. If you do not want the interfaces to remain online during the switchover for the other FEB, configure a primary FEB for the redundancy group.

Failover to a backup FEB occurs automatically if a FEB in a redundancy group fails. You can disable automatic failover for any redundancy group by including the **no-auto-failover** statement at the **[edit chassis redundancy feb redundancy-group group-name]** hierarchy level.

You can also initiate a manual switchover by issuing the **request chassis redundancy feb slot slot-number switch-to-backup** command, where **slot-number** is the number of the active FEB. For more information, see the [CLI Explorer](#).

The following conditions result in failover as long as the backup FEB in a redundancy group is available:

- The FEB is absent.
- The FEB experienced a hard error while coming online.
- A software failure on the FEB resulted in a crash.
- Ethernet connectivity from a FEB to a Routing Engine failed.
- A hard error on the FEB, such as a power failure, occurred.
- The FEB was disabled when the offline button for the FEB was pressed.
- The software watchdog timer on the FEB expired.
- Errors occurred on the links between all the active fabric planes and the FEB. This situation results in failover to the backup FEB if it has at least one valid fabric link.
- Errors occurred on the link between the FEB and all of the FPCs connected to it.

After a switchover occurs, a backup FEB is no longer available for the redundancy group. You can revert from the backup FEB to the previously active FEB by issuing the operational mode command **request chassis redundancy feb slot slot-number revert-from-backup**, where **slot-number** is the number of the previously active FEB. For more information, see the [CLI Explorer](#).

When you revert from the backup FEB, it becomes available again for a switchover. If the redundancy group does not have a primary FEB, the backup FEB reboots after you revert back to the previously active FEB. If the FEB to which you revert back is not a primary FEB, the backup FEB is rebooted so that it can align with the state of the primary FEB.

If you modify the configuration for an existing redundancy group so that a FEB connects to a different FPC, the FEB is rebooted unless the FEB was already connected to one or two Type 1 FPCs and the change only resulted in the FEB being connected either to one additional or one fewer Type 1 FPC. For more information about how to map a connection between an FPC and a FEB, see the *Junos OS Administration Library*. If you change the primary FEB in a redundancy group, the backup FEB is rebooted. The FEB is also rebooted

if you change a backup FEB to a nonbackup FEB or change an active FEB to a backup FEB.

To view the status of configured FEB redundancy groups, issue the **show chassis redundancy feb** operational mode command. For more information, see the [CLI Explorer](#).

## Redundant SSBs on the M20 Router

The System and Switch Board (SSB) on the M20 router performs the following major functions:

- Shared memory management on the FPCs—The Distributed Buffer Manager ASIC on the SSB uniformly allocates incoming data packets throughout shared memory on the FPCs.
- Outgoing data cell transfer to the FPCs—A second Distributed Buffer Manager ASIC on the SSB passes data cells to the FPCs for packet reassembly when the data is ready to be transmitted.
- Route lookups—The Internet Processor ASIC on the SSB performs route lookups using the forwarding table stored in SSRAM. After performing the lookup, the Internet Processor ASIC informs the midplane of the forwarding decision, and the midplane forwards the decision to the appropriate outgoing interface.
- System component monitoring—The SSB monitors other system components for failure and alarm conditions. It collects statistics from all sensors in the system and relays them to the Routing Engine, which sets the appropriate alarm. For example, if a temperature sensor exceeds the first internally defined threshold, the Routing Engine issues a “high temp” alarm. If the sensor exceeds the second threshold, the Routing Engine initiates a system shutdown.
- Exception and control packet transfer—The Internet Processor ASIC passes exception packets to a microprocessor on the SSB, which processes almost all of them. The remaining packets are sent to the Routing Engine for further processing. Any errors that originate in the Packet Forwarding Engine and are detected by the SSB are sent to the Routing Engine using system log messages.
- FPC reset control—The SSB monitors the operation of the FPCs. If it detects errors in an FPC, the SSB attempts to reset the FPC. After three unsuccessful resets, the SSB takes the FPC offline and informs the Routing Engine. Other FPCs are unaffected, and normal system operation continues.

The M20 router holds up to two SSBs. One SSB is configured to act as the master and the other is configured to serve as a backup in case the master fails. You can initiate a manual switchover by issuing the **request chassis ssb master switch** command. For more information, see the [CLI Explorer](#).

## Redundant SFMs on the M40e and M160 Routers

The M40e and M160 routers have redundant Switching and Forwarding Modules (SFMs). The SFMs contain the Internet Processor II ASIC and two Distributed Buffer Manager ASICs. SFMs ensure that all traffic leaving the FPCs is handled properly. SFMs provide route lookup, filtering, and switching.

The M40e router holds up to two SFMs, one that is configured to act as the master and the other configured to serve as a backup in case the master fails. Removing the standby SFM has no effect on router function. If the active SFM fails or is removed from the chassis, forwarding halts until the standby SFM boots and becomes active. It takes approximately 1 minute for the new SFM to become active. Synchronizing router configuration information can take additional time, depending on the complexity of the configuration.

The M160 router holds up to four SFMs. All SFMs are active at the same time. A failure or taking an SFM offline has no effect on router function. Forwarding continues uninterrupted.

You can initiate a manual switchover by issuing the **request chassis sfm master switch** command. For more information, see the [CLI Explorer](#).

#### Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 17](#)
- [Understanding Routing Engine Redundancy on Juniper Networks Routers](#)
- [Configuring CFEB Redundancy on the M10i Router on page 29](#)
- [Configuring FEB Redundancy on the M120 Router on page 30](#)
- [Configuring SFM Redundancy on M40e and M160 Routers on page 33](#)
- [Configuring SSB Redundancy on the M20 Router on page 34](#)
- *show chassis redundancy feb*
- *request chassis cb*

## Configuring CFEB Redundancy on the M10i Router

The Compact Forwarding Engine Board (CFEB) on the M10i router provides route lookup, filtering, and switching on incoming data packets, and then directs outbound packets to the appropriate interface for transmission to the network. The CFEB communicates with the Routing Engine using a dedicated 100-Mbps Fast Ethernet link that transfers routing table data from the Routing Engine to the forwarding table in the integrated ASIC. The link is also used to transfer from the CFEB to the Routing Engine routing link-state updates and other packets destined for the router that have been received through the router interfaces.

To configure a CFEB redundancy group, include the following statements at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
  cfeb slot-number (always | preferred);
```

*slot-number* can be 0 or 1.

**always** defines the CFEB as the sole device.

**preferred** defines a preferred CFEB.

To manually switch CFEB mastership, issue the **request chassis cfef master switch** command. To view CFEB status, issue the **show chassis cfef** command.

**Related Documentation**

- [Understanding Switching Control Board Redundancy on page 25](#)

---

## Configuring Switching Control Board Redundancy

---

- [Configuring CFEB Redundancy on the M10i Router on page 30](#)
- [Configuring FEB Redundancy on the M120 Router on page 30](#)
- [Example: Configuring FEB Redundancy on M120 Routers on page 31](#)
- [Configuring SFM Redundancy on M40e and M160 Routers on page 33](#)
- [Configuring SSB Redundancy on the M20 Router on page 34](#)

### Configuring CFEB Redundancy on the M10i Router

The Compact Forwarding Engine Board (CFEB) on the M10i router provides route lookup, filtering, and switching on incoming data packets, and then directs outbound packets to the appropriate interface for transmission to the network. The CFEB communicates with the Routing Engine using a dedicated 100-Mbps Fast Ethernet link that transfers routing table data from the Routing Engine to the forwarding table in the integrated ASIC. The link is also used to transfer from the CFEB to the Routing Engine routing link-state updates and other packets destined for the router that have been received through the router interfaces.

To configure a CFEB redundancy group, include the following statements at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
  cfef slot-number (always | preferred);
```

*slot-number* can be 0 or 1.

**always** defines the CFEB as the sole device.

**preferred** defines a preferred CFEB.

To manually switch CFEB mastership, issue the **request chassis cfef master switch** command. To view CFEB status, issue the **show chassis cfef** command.

**See Also**

- [Understanding Switching Control Board Redundancy on page 25](#)

### Configuring FEB Redundancy on the M120 Router

To configure a FEB redundancy group for the M120 router, include the following statements at the **[edit chassis redundancy feb]** hierarchy level:

```
[edit chassis redundancy feb]
  redundancy-group group-name {
    description description;
    feb slot-number (backup | primary);
```

```
no-auto-failover;
}
```

**group-name** is the unique name for the redundancy group. The maximum length is 39 alphanumeric characters.

**slot-number** is the slot number of each FEB you want to include in the redundancy group. The range is from 0 through 5. You must specify exactly one FEB as a backup FEB per redundancy group. Include the **backup** keyword when configuring the backup FEB and make sure that the FEB is not connected to an FPC.

Include the **primary** keyword to optionally specify one primary FEB per redundancy group. When the **primary** keyword is specified for a particular FEB, that FEB is configured for 1:1 redundancy. With 1:1 redundancy, the backup FEB contains the same forwarding state as the primary FEB. When no FEB in the redundancy group is configured as a primary FEB, the redundancy group is configured for *n*:1 redundancy. In this case, the backup FEB has no forwarding state. When a FEB fails, the forwarding state must be downloaded from the Routing Engine to the backup FEB before forwarding continues.

A combination of 1:1 and *n*:1 redundancy is possible when more than two FEBs are present in a group. The backup FEB contains the same forwarding state as the primary FEB, so that when the primary FEB fails, 1:1 failover is in effect. When a nonprimary FEB fails, the backup FEB must be rebooted so that the forwarding state from the nonprimary FEB is installed on the backup FEB before it can continue forwarding.

You can optionally include the **description** statement to describe a redundancy group.

Automatic failover is enabled by default. To disable automatic failover, include the **no-auto-failover** statement. If you disable automatic failover, you can perform only a manual switchover using the operational command **request chassis redundancy feb slot slot-number switch-to-backup**.

To view FEB status, issue the **show chassis feb** command. For more information, see the [CLI Explorer](#).

- See Also**
- [Understanding Switching Control Board Redundancy on page 25](#)
  - [Example: Configuring FEB Redundancy on M120 Routers on page 31](#)

### Example: Configuring FEB Redundancy on M120 Routers

In the following configuration, two FEB redundancy groups are created:

- A FEB redundancy group named **group0** with the following properties:
  - Contains three FEBs (0 through 2).
  - Has a primary FEB (2).
  - Has a unique backup FEB (0).
  - Automatic failover is disabled.

When an active FEB in **group0** fails, automatic failover to the backup FEB does not occur. For **group0**, you can only perform a manual switchover.



- A FEB redundancy group named **group1** with the following properties:
  - Two FEBs (3 and 5). There is no primary FEB.
  - A unique backup FEB (5).
  - Automatic failover is enabled by default.

When **feb 3** in **group1** fails, an automatic failover occurs.

Because you must explicitly configure an FPC *not* to connect to the backup FEB, connectivity is set to none between **fpc 0** and **feb 0** and between **fpc 5** and **feb 5**.



**NOTE:** For information about the **fpc-feb-connectivity** statement, see the *Junos OS Administration Library*.

FPC to primary FEB connectivity is not explicitly configured, so by default, the software automatically assigns connectivity based on the numerical order of the FPCs.

```
[edit]
chassis {
  fpc-feb-connectivity {
    fpc 0 feb none;
    fpc 5 feb none;
  }
  redundancy feb {
    redundancy-group group0 {
      description "Interfaces to Customer X";
      feb 2 primary;
      feb 1;
      feb 0 backup;
      no-auto-failover;
    }
    redundancy-group group1 {
      feb 3;
      feb 5 backup;
    }
  }
}
```

- See Also**
- [Understanding Switching Control Board Redundancy on page 25](#)
  - [Configuring FEB Redundancy on the M120 Router on page 30](#)

## Configuring SFM Redundancy on M40e and M160 Routers

By default, the Switching and Forwarding Module (SFM) in slot 0 is the master and the SFM in slot 1 is the backup. To modify the default configuration, include the **sfm** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
sfm slot-number (always | preferred);
```

On the M40e router, **slot-number** is 0 or 1. On the M160 router, **slot-number** is 0 through 3.

**always** defines the SFM as the sole device.

**preferred** defines a preferred SFM.

To manually switch mastership between SFMs, issue the **request chassis sfm master switch** command. To view SFM status, issue the **show chassis sfm** command. For more information, see the [CLI Explorer](#).

**See Also** • [Understanding Switching Control Board Redundancy on page 25](#)

## Configuring SSB Redundancy on the M20 Router

For M20 routers with two System and Switch Boards (SSBs), you can configure which SSB is the master and which is the backup. By default, the SSB in slot 0 is the master and the SSB in slot 1 is the backup. To modify the default configuration, include the **ssb** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
  ssb slot-number (always | preferred);
```

*slot-number* is 0 or 1.

**always** defines the SSB as the sole device.

**preferred** defines a preferred SSB.

To manually switch mastership between SSBs, issue the **request chassis ssb master switch** command.

To display SSB status information, issue the **show chassis ssb** command. The command output displays the number of times the mastership has changed, the SSB slot number, and the current state of the SSB: master, backup, or empty. For more information, see the [CLI Explorer](#).

**See Also** • [Understanding Switching Control Board Redundancy on page 25](#)

### Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 17](#)
- [Understanding Routing Engine Redundancy on Juniper Networks Routers](#)
- *show chassis redundancy feb*
- *request chassis cb*

---

## Configuring SFM Redundancy on M40e and M160 Routers

By default, the Switching and Forwarding Module (SFM) in slot 0 is the master and the SFM in slot 1 is the backup. To modify the default configuration, include the **sfm** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
  sfm slot-number (always | preferred);
```

On the M40e router, **slot-number** is 0 or 1. On the M160 router, **slot-number** is 0 through 3.

**always** defines the SFM as the sole device.

**preferred** defines a preferred SFM.

To manually switch mastership between SFMs, issue the **request chassis sfm master switch** command. To view SFM status, issue the **show chassis sfm** command. For more information, see the [CLI Explorer](#).

**Related Documentation** • [Understanding Switching Control Board Redundancy on page 25](#)

---

## Configuring SSB Redundancy on the M20 Router

---

For M20 routers with two System and Switch Boards (SSBs), you can configure which SSB is the master and which is the backup. By default, the SSB in slot 0 is the master and the SSB in slot 1 is the backup. To modify the default configuration, include the **ssb** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
  ssb slot-number (always | preferred);
```

**slot-number** is 0 or 1.

**always** defines the SSB as the sole device.

**preferred** defines a preferred SSB.

To manually switch mastership between SSBs, issue the **request chassis ssb master switch** command.

To display SSB status information, issue the **show chassis ssb** command. The command output displays the number of times the mastership has changed, the SSB slot number, and the current state of the SSB: master, backup, or empty. For more information, see the [CLI Explorer](#).

**Related Documentation** • [Understanding Switching Control Board Redundancy on page 25](#)



## CHAPTER 3

# Configuring Bidirectional Forwarding Detection (BFD)

- [Configuring BFD for Static Routes on page 37](#)
- [Configuring BFD for BGP on page 48](#)
- [Configuring BFD for OSPF on page 59](#)
- [Configuring BFD for IS-IS on page 65](#)
- [Configuring BFD for RIP on page 74](#)
- [Configuring BFD for PIM on page 81](#)
- [Configuring Independent Micro BFD Sessions for LAG on page 83](#)
- [Understanding Distributed BFD on page 100](#)

## Configuring BFD for Static Routes

---

- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 37](#)
- [Understanding Static Route State When BFD is in Admin Down State on page 42](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 42](#)

## Understanding BFD for Static Routes for Faster Network Failure Detection

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the static route failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The

transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

By default, BFD is supported on single-hop static routes.

To enable failure detection, include the **bfd-liveness-detection** statement in the static route configuration.



**NOTE:** Starting with Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, the **bfd-liveness-detection** command includes the description field. The description is an attribute under the **bfd-liveness-detection** object and it is supported only on SRX Series devices. This field is applicable only for the static routes.

In Junos OS Release 9.1 and later, the BFD protocol is supported for IPv6 static routes. Global unicast and link-local IPv6 addresses are supported for static routes. The BFD protocol is not supported on multicast or anycast IPv6 addresses. For IPv6, the BFD protocol supports only static routes and only in Junos OS Release 9.3 and later. IPv6 for BFD is also supported for the eBGP protocol.



**NOTE:** Inline BFD is supported on PTX5000 routers with third-generation FPCs starting in Junos OS Release 15.1F3 and 16.1R2. Inline BFD is supported on PTX3000 routers with third-generation FPCs starting in Junos OS Release 15.1F6 and 16.1R2.

There are three types of BFD sessions based on the source from which BFD packets are sent to the neighbors. Different types of BFD sessions and their descriptions are:

Type of BFD session	Description
Non-distributed BFD	BFD sessions running completely on the Routing Engine.
Distributed BFD	BFD sessions running completely on the Packet Forwarding Engine.
Inline BFD	BFD sessions running on the FPC hardware.
<p><b>NOTE:</b> Starting in Junos OS Release 13.3, inline BFD is supported only on static MX Series routers with MPCs/MICs that have configured <b>enhanced-ip</b>.</p> <p><b>NOTE:</b> Starting in Junos OS Release 16.1R1, the inline BFD sessions are supported on integrated routing and bridging (IRB) interfaces.</p>	

To configure the BFD protocol for IPv6 static routes, include the **bfd-liveness-detection** statement at the **[edit routing-options rib inet6.0 static route destination-prefix]** hierarchy level.

In Junos OS Release 8.5 and later, you can configure a hold-down interval to specify how long the BFD session must remain up before a state change notification is sent.

To specify the hold-down interval, include the **holddown-interval** statement in the BFD configuration.

You can configure a number in the range from 0 through 255,000 milliseconds. The default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.



**NOTE:** If a single BFD session includes multiple static routes, the hold-down interval with the highest value is used.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement in the BFD configuration.

This value represents both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



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

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 (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

To specify the minimum receive interval for failure detection, include the **minimum-receive-interval** statement in the BFD configuration. This value represents the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down, include the **multiplier** statement in the BFD configuration.

The default value is 3. You can configure a number in the range from 1 through 255.

To specify a threshold for detecting the adaptation of the detection time, include the **threshold** statement in the BFD configuration.

When the BFD session detection time adapts to a value equal to or higher than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the **minimum-interval** or the **minimum-receive-interval** value. The threshold must be a higher value than the multiplier for either of these configured values. For example if the **minimum-receive-interval** is 300 ms and the **multiplier** is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value higher than 900.

To specify the minimum transmit interval for failure detection, include the **transmit-interval** **minimum-interval** statement in the BFD configuration.

This value represents the minimum interval after which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the threshold for the adaptation of the transmit interval, include the **transmit-interval threshold** statement in the BFD configuration.

The threshold value must be greater than the transmit interval. When the BFD session transmit time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the value for the **minimum-interval** or the **minimum-receive-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level. The threshold must be a higher value than the multiplier for either of these configured values.

To specify the BFD version, include the **version** statement in the BFD configuration. The default is to have the version detected automatically.

To include an IP address for the next hop of the BFD session, include the **neighbor** statement in the BFD configuration.





**NOTE:** You must configure the `neighbor` statement if the next hop specified is an interface name. If you specify an IP address as the next hop, that address is used as the neighbor address for the BFD session.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement in the BFD configuration.



**NOTE:** We recommend that you not disable BFD adaptation unless it is preferable *not* to have BFD adaptation in your network.



**NOTE:** If BFD is configured only on one end of a static route, the route is removed from the routing table. BFD establishes a session when BFD is configured on both ends of the static route.

BFD is not supported on ISO address families in static routes. BFD does support IS-IS.

If you configure graceful Routing Engine switchover (GRES) at the same time as BFD, GRES does not preserve the BFD state information during a failover.

- See Also**
- *Enabling Dedicated and Real-Time BFD*
  - [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 42](#)
  - *Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection*

## Understanding Static Route State When BFD is in Admin Down State

The Bidirectional Forwarding Detection (BFD) Admin Down state is used to bring down a BFD session administratively (applicable for normal BFD session and micro BFD session), to protect client applications from BFD configuration removal, license issues, and clearing of BFD sessions.

When BFD enters the Admin Down state, BFD notifies the new state to its peer for a failure detection time and after the time expires, the client stops transmitting packets.

For the Admin Down state to work, the peer, which receives the Admin Down state notification, must have the capability to distinguish between administratively down state and real link failure.

A BFD session moves to the Admin Down state under the following conditions:

- If BFD configuration is removed for the last client tied to a BFD session, BFD moves to Admin Down state and communicates the change to the peer, to enable the client protocols without going down.
- If BFD license is removed on the client, BFD moves to Admin Down state and communicates the change to the remote system to enable the client protocols without going down.
- When **clear bfd session** command is executed, the BFD sessions move to Admin Down state before restarting. This **clear bfd session** command also ensures that the client applications are not impacted.

Starting from Junos OS 16.1R1 release, you can set the state of static route in BFD Admin Down state by configuring one of the following commands:

- **set routing-options static static-route bfd-admin-down active**—BFD Admin Down state pulls down the static route.
- **set routing-options static static-route bfd-admin-down passive**—BFD Admin Down state does not pull down the static route.

- See Also**
- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 37](#)
  - [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 42](#)

## Example: Configuring BFD for Static Routes for Faster Network Failure Detection

This example shows how to configure Bidirectional Forwarding Detection (BFD) for static routes.

- [Requirements on page 43](#)
- [Overview on page 43](#)
- [Configuration on page 43](#)
- [Verification on page 47](#)

## Requirements

In this example, no special configuration beyond device initialization is required.

## Overview


There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

Figure 1 on page 43 shows the sample network.

*Figure 1: Customer Routes Connected to a Service Provider*



## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device B**     **set interfaces ge-1/2/0 unit 0 description B->D**

```
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description
  Site-xxx
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
```

**Device D**

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

```
[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```

2. On Device B, create a static route and set the next-hop address.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
```

3. On Device B, configure BFD for the static route.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description
  Site-xxx
```

4. On Device B, configure tracing operations for BFD.

```
[edit protocols]
user@B# set bfd traceoptions file bfd-trace
user@B# set bfd traceoptions flag all
```

5. If you are done configuring Device B, commit the configuration.

```
[edit]
user@B# commit
```

6. On Device D, configure the interfaces.

```
[edit interfaces]
user@D# set ge-1/2/0 unit 1 description D->B
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
user@D# set lo0 unit 2 family inet address 192.168.47.6/32
```

7. On Device D, create a static route and set the next-hop address.

```
[edit routing-options]
user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
```

8. On Device D, configure BFD for the static route.

```
[edit routing-options]
user@D# set static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
```

9. On Device D, configure tracing operations for BFD.

```
[edit protocols]
user@D# set bfd traceoptions file bfd-trace
user@D# set bfd traceoptions flag all
```

10. If you are done configuring Device D, commit the configuration.

```
[edit]
user@D# commit
```

### Results

Confirm your configuration by issuing the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device B user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}
lo0 {
  unit 57 {
    family inet {
      address 10.0.0.1/32;
    }
  }
}
```

```
        address 10.0.0.2/32;
    }
}
```

```
user@D# show protocols
bfd {
  traceoptions {
    file bfd-trace;
    flag all;
  }
}
```

```
user@B# show routing-options
static {
  route 192.168.47.0/24 {
    next-hop 172.16.1.2;
    bfd-liveness-detection {
      description Site- xxx;
      minimum-interval 1000;
    }
  }
}
```

**Device D**

```
user@D# show interfaces
ge-1/2/0 {
  unit 1 {
    description D->B;
    family inet {
      address 172.16.1.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.47.5/32;
      address 192.168.47.6/32;
    }
  }
}
```

```
user@D# show routing-options
static {
  route 0.0.0.0/0 {
    next-hop 172.16.1.1;
    bfd-liveness-detection {
      description Site - xxx;
      minimum-interval 1000;
    }
  }
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 47](#)
- [Viewing Detailed BFD Events on page 48](#)

### Verifying That BFD Sessions Are Up

**Purpose** Verify that the BFD sessions are up, and view details about the BFD sessions.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	lt-1/2/0.0	3.000	1.000	3

Client Static, description Site-xxx, TX interval 1.000, RX interval 1.000  
Session up time 00:14:30  
Local diagnostic None, remote diagnostic None  
Remote state Up, version 1  
Replicated, routing table index 172  
Min async interval 1.000, min slow interval 1.000  
Adaptive async TX interval 1.000, RX interval 1.000  
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3  
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3  
Local discriminator 2, remote discriminator 1  
Echo mode disabled/inactive

1 sessions, 1 clients  
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps



**NOTE:** The description Site- <xxx> is supported only on the SRX Series devices.

If each client has more than one description field, then it displays "and more" along with the first description field.

```
user@D> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.1	Up	lt-1/2/0.1	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000  
Session up time 00:14:35  
Local diagnostic None, remote diagnostic None  
Remote state Up, version 1  
Replicated, routing table index 170  
Min async interval 1.000, min slow interval 1.000  
Adaptive async TX interval 1.000, RX interval 1.000  
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3  
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3  
Local discriminator 1, remote discriminator 2

Echo mode disabled/inactive

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

**Meaning** The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

#### *Viewing Detailed BFD Events*

**Purpose** View the contents of the BFD trace file to assist in troubleshooting, if needed.

**Action** From operational mode, enter the **file show /var/log/bfd-trace** command.

```
user@B> file show /var/log/bfd-trace
Nov 23 14:26:55 Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 72
Nov 23 14:26:55 PPM Trace: BFD periodic xmit rt tbl index 172
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 108:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 83: (hex) 70 70 6d 64 5f 62 66 64 5f 73 65 6e 64
6d 73 67 20 3a 20
Nov 23 14:26:55 PPM Trace: ppm_bfd_sendmsg : socket 12 len 24, ifl 78 src
172.16.1.1 dst 172.16.1.2 errno 65
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 93:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 68: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 74
```

**Meaning** BFD messages are being written to the trace file.

**See Also** • [Understanding BFD for Static Routes for Faster Network Failure Detection on page 37](#)

---

## Configuring BFD for BGP

- [Understanding BFD for BGP on page 48](#)
- [Example: Configuring BFD on Internal BGP Peer Sessions on page 50](#)

### Understanding BFD for BGP

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The



failure detection timers for BFD have shorter time limits than default failure detection mechanisms for BGP, so they provide faster detection.



**NOTE:** Configuring both BFD and graceful restart for BGP on the same device is counterproductive. When an interface goes down, BFD detects this instantly, stopping traffic forwarding and the BGP session goes down conversely graceful restart forwards traffic despite the interface failure; this behavior might cause network issues. Hence we do not recommend configuring both BFD and graceful restart on the same device.

The BFD failure detection timers can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



**NOTE:** On all SRX Series devices, high CPU utilization triggered for reasons such as CPU intensive commands and SNMP walks causes the BFD protocol to flap while processing large BGP updates. (Platform support depends on the Junos OS release in your installation.)

Starting with Junos OS Release 15.1X49-D100, SRX340, SRX345, and SRX1500 devices support dedicated BFD.

Starting with Junos OS Release 15.1X49-D100, SRX300 and SRX320 devices support real-time BFD.

Starting with Junos OS Release 15.1X49-D110, SRX550M devices support dedicated BFD.

Starting with Junos OS Release 12.3X48-D60, SRX240, SRX550, and SRX650 devices support dedicated BFD.

Starting with Junos OS Release 12.3X48-D60, SRX100, SRX110, SRX210, and SRX220 devices support real-time BFD.

In Junos OS Release 8.3 and later, BFD is supported on internal BGP (IBGP) and multihop external BGP (EBGP) sessions as well as on single-hop EBGP sessions. In Junos OS Release 9.1 through Junos OS Release 11.1, BFD supports IPv6 interfaces in static routes only. In Junos OS Release 11.2 and later, BFD supports IPv6 interfaces with BGP.

- See Also**
- [Enabling Dedicated and Real-Time BFD](#)
  - [Example: Configuring BFD on Internal BGP Peer Sessions on page 50](#)

## Example: Configuring BFD on Internal BGP Peer Sessions

This example shows how to configure internal BGP (IBGP) peer sessions with the Bidirectional Forwarding Detection (BFD) protocol to detect failures in a network.

- [Requirements on page 50](#)
- [Overview on page 50](#)
- [Configuration on page 52](#)
- [Verification on page 56](#)

### Requirements

---

No special configuration beyond device initialization is required before you configure this example.

### Overview

---

The minimum configuration to enable BFD on IBGP sessions is to include the **bfd-liveness-detection minimum-interval** statement in the BGP configuration of all neighbors participating in the BFD session. The **minimum-interval** statement specifies the minimum transmit and receive intervals for failure detection. Specifically, this value represents the minimum interval after which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.

Optionally, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements. For information about these and other optional BFD configuration statements, see **bfd-liveness-detection**.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 milliseconds for Routing Engine-based sessions and less than 10 milliseconds for distributed BFD sessions can cause undesired BFD flapping.


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

- To prevent BFD flapping during the general Routing Engine switchover event, specify a minimum interval of 5000 milliseconds for Routing Engine-based sessions. This minimum value is required because, during the general Routing Engine switchover event, processes such as RPD, MIBD, and SNMPD utilize CPU resources for more than the specified threshold value. Hence, BFD processing and scheduling is affected because of this lack of CPU resources.
- For BFD sessions to remain up during the dual chassis cluster control link scenario, when the first control link fails, specify the minimum interval of 6000 milliseconds to prevent the LACP from flapping on the secondary node for Routing Engine-based sessions.
- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 milliseconds for Routing Engine-based sessions and 100 milliseconds 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 (NSR) is configured, specify a minimum interval of 2500 milliseconds for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

BFD is supported on the default routing instance (the main router), routing instances, and logical systems. This example shows BFD on logical systems.

Figure 2 on page 52 shows a typical network with internal peer sessions.

Figure 2: Typical Network with IBGP Sessions



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Device A

```
set logical-systems A interfaces lt-1/2/0 unit 1 description to-B
set logical-systems A interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems A interfaces lt-1/2/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-bfd
set logical-systems A protocols bgp group internal-peers traceoptions flag bfd detail
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers bfd-liveness-detection
    minimum-interval 1000
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-1/2/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
set logical-systems A routing-options autonomous-system 17
```

#### Device B

```
set logical-systems B interfaces lt-1/2/0 unit 2 description to-A
set logical-systems B interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-1/2/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-1/2/0 unit 5 description to-C
set logical-systems B interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-1/2/0 unit 5 family inet address 10.10.10.5/30
```

```

set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers bfd-liveness-detection
    minimum-interval 1000
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17

```

**Device C**

```

set logical-systems C interfaces lt-1/2/0 unit 6 description to-B
set logical-systems C interfaces lt-1/2/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-1/2/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-1/2/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers bfd-liveness-detection
    minimum-interval 1000
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-1/2/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17

```

### Configuring Device A

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device A:

1. Set the CLI to Logical System A.  

```
user@host> set cli logical-system A
```
2. Configure the interfaces.  

```
[edit interfaces lt-1/2/0 unit 1]
user@host:A# set description to-B
user@host:A# set encapsulation ethernet
user@host:A# set peer-unit 2
```

```
user@host:A# set family inet address 10.10.10.1/30
```

```
[edit interfaces lo0 unit 1]
```

```
user@host:A# set family inet address 192.168.6.5/32
```

3. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
```

```
user@host:A# set type internal
```

```
user@host:A# set local-address 192.168.6.5
```

```
user@host:A# set export send-direct
```

```
user@host:A# set neighbor 192.163.6.4
```

```
user@host:A# set neighbor 192.168.40.4
```

4. Configure BFD.

```
[edit protocols bgp group internal-peers]
```

```
user@host:A# set bfd-liveness-detection minimum-interval 1000
```

You must configure the same minimum interval on the connecting peer.

5. (Optional) Configure BFD tracing.

```
[edit protocols bgp group internal-peers]
```

```
user@host:A# set traceoptions file bgp-bfd
```

```
user@host:A# set traceoptions flag bfd detail
```

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@host:A# set interface lo0.1 passive
```

```
user@host:A# set interface lt-1/2/0.1
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
```

```
user@host:A# set from protocol direct
```

```
user@host:A# set then accept
```

8. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
```

```
user@host:A# set router-id 192.168.6.5
```

```
user@host:A# set autonomous-system 17
```

9. If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps to configure Device B and Device C.

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host:A# show interfaces
lt-1/2/0 {
  unit 1 {
    description to-B;
    encapsulation ethernet;
    peer-unit 2;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@host:A# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@host:A# show protocols
bgp {
  group internal-peers {
    type internal;
    traceoptions {
      file bgp-bfd;
      flag bfd detail;
    }
    local-address 192.168.6.5;
    export send-direct;
    bfd-liveness-detection {
      minimum-interval 1000;
    }
    neighbor 192.163.6.4;
    neighbor 192.168.40.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface lt-1/2/0.1;
  }
}

```

```

}

user@host:A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;

```

## Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Is Enabled on page 56](#)
- [Verifying That BFD Sessions Are Up on page 56](#)
- [Viewing Detailed BFD Events on page 57](#)
- [Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface on page 58](#)

### Verifying That BFD Is Enabled

**Purpose** Verify that BFD is enabled between the IBGP peers.

**Action** From operational mode, enter the **show bgp neighbor** command. You can use the **| match bfd** filter to narrow the output.

```

user@host:A> show bgp neighbor | match bfd
Options: <BfdEnabled>
  BFD: enabled, up
  Trace file: /var/log/A/bgp-bfd size 131072 files 10
Options: <BfdEnabled>
  BFD: enabled, up
  Trace file: /var/log/A/bgp-bfd size 131072 files 10

```

**Meaning** The output shows that Logical System A has two neighbors with BFD enabled. When BFD is not enabled, the output displays **BFD: disabled, down**, and the **<BfdEnabled>** option is absent. If BFD is enabled and the session is down, the output displays **BFD: enabled, down**. The output also shows that BFD-related events are being written to a log file because trace operations are configured.

### Verifying That BFD Sessions Are Up

**Purpose** Verify that the BFD sessions are up, and view details about the BFD sessions.

**Action** From operational mode, enter the **show bfd session extensive** command.

```

user@host:A> show bfd session extensive

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.163.6.4	Up		3.000	1.000	3

```

Client BGP, TX interval 1.000, RX interval 1.000
Session up time 00:54:40
Local diagnostic None, remote diagnostic None
Remote state Up, version 1

```



```

Logical system 12, routing table index 25
Min async interval 1.000, min slow interval 1.000
Adaptive async TX interval 1.000, RX interval 1.000
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
Local discriminator 10, remote discriminator 9
Echo mode disabled/inactive
Multi-hop route table 25, local-address 192.168.6.5

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.168.40.4	Up		3.000	1.000	3

```

Client BGP, TX interval 1.000, RX interval 1.000
Session up time 00:48:03
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Logical system 12, routing table index 25
Min async interval 1.000, min slow interval 1.000
Adaptive async TX interval 1.000, RX interval 1.000
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
Local discriminator 14, remote discriminator 13
Echo mode disabled/inactive
Multi-hop route table 25, local-address 192.168.6.5

```

```

2 sessions, 2 clients
Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps

```

**Meaning** The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

### Viewing Detailed BFD Events

**Purpose** View the contents of the BFD trace file to assist in troubleshooting, if needed.

**Action** From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```

user@host:A> file show /var/log/A/bgp-bfd
Aug 15 17:07:25 trace_on: Tracing to "/var/log/A/bgp-bfd" started
Aug 15 17:07:26.492190 bgp_peer_init: BGP peer 192.163.6.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:26.493176 bgp_peer_init: BGP peer 192.168.40.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:32.597979 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:07:32.599623 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:07:36.869394 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:07:36.870624 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:04.599220 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:08:04.601135 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):

```

```

No route to host
Aug 15 17:08:08.869717 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:08:08.869934 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
  No route to host
Aug 15 17:08:36.603544 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:08:36.606726 bgp_read_message: 192.163.6.4 (Internal AS 17): 0 bytes
buffered
Aug 15 17:08:36.609119 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:36.734033 advertising receiving-speaker only capability to neighbor
192.168.40.4 (Internal AS 17)
Aug 15 17:08:36.738436 Initiated BFD session to peer 192.168.40.4 (Internal AS
17): address=192.168.40.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:40.537552 BFD session to peer 192.163.6.4 (Internal AS 17) up
Aug 15 17:08:40.694410 BFD session to peer 192.168.40.4 (Internal AS 17) up

```

**Meaning** Before the routes are established, the **No route to host** message appears in the output. After the routes are established, the last two lines show that both BFD sessions come up.

#### *Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface*

**Purpose** Check to see what happens after bringing down a router or switch and then bringing it back up. To simulate bringing down a router or switch, deactivate the loopback interface on Logical System B.

**Action** 1. From configuration mode, enter the **deactivate logical-systems B interfaces lo0 unit 2 family inet** command.

```

user@host:A# deactivate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit

```

2. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```

user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:20:55.995648 bgp_read_v4_message:9747: NOTIFICATION received from
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 6 (Other Configuration
Change)
Aug 15 17:20:56.004508 Terminated BFD session to peer 192.163.6.4 (Internal
AS 17)
Aug 15 17:21:28.007755 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:21:28.008597 bgp_connect_start: connect 192.163.6.4 (Internal AS
17): No route to host

```

3. From configuration mode, enter the **activate logical-systems B interfaces lo0 unit 2 family inet** command.

```

user@host:A# activate logical-systems B interfaces lo0 unit 2 family inet

```

```
user@host:A# commit
```

4. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```
user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:25:53.623743 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:25:53.631314 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:25:57.570932 BFD session to peer 192.163.6.4 (Internal AS 17) up
```

**See Also** • [Understanding BFD Authentication for BGP](#)

## Configuring BFD for OSPF

- [Understanding BFD for OSPF on page 59](#)
- [Example: Configuring BFD for OSPF on page 61](#)

### Understanding BFD for OSPF

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchange BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the OSPF failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



**NOTE:** BFD is supported for OSPFv3 in Junos OS Release 9.3 and later.

You can configure the following BFD protocol settings:

- **detection-time threshold**—Threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the configured threshold, a single trap and a single system log message are sent.
- **full-neighbors-only**—Ability to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors. This setting is available in Junos OS Release 9.5 and later.
- **minimum-interval**—Minimum transmit and receive interval for failure detection. This setting configures both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. Both intervals are in milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.



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

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 no less than 500 ms. An interval of 1000 ms is recommended to avoid any instability issues.
  - 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 (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. Without NSR, Routing Engine-based sessions can have a minimum interval of 100 ms. In OSPFv3, BFD is always based in the Routing Engine, meaning that BFD is not distributed. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
  - On a single QFX5100 switch, when you add a QFX-EM-4Q expansion module, specify a minimum interval higher than 1000 ms.
- 
- **minimum-receive-interval**—Minimum receive interval for failure detection. This setting configures the minimum receive interval, in milliseconds, after which the routing device expects to receive a hello packet from a neighbor with which it has established a BFD session. You can also specify the minimum receive interval using the **minimum-interval** statement.

- **multiplier**—Multiplier for hello packets. This setting configures the number of hello packets that are not received by a neighbor, which causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down.
- **no-adaptation**—Disables BFD adaptation. This setting disables BFD sessions from adapting to changing network conditions. This setting is available in Junos OS Release 9.0 and later.



**NOTE:** We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

- **transmit-interval minimum-interval**—Minimum transmit interval for failure detection. This setting configures the minimum transmit interval, in milliseconds, at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can also specify the minimum transmit interval using the **minimum-interval** statement.
- **transmit-interval threshold**—Threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The threshold value must be greater than the minimum transmit interval. If you attempt to commit a configuration with a threshold value less than the minimum transmit interval, the routing device displays an error and does not accept the configuration.
- **version**—BFD version. This setting configures the BFD version used for detection. You can explicitly configure BFD version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version automatically, which is either 0 or 1.

You can also trace BFD operations for troubleshooting purposes.

- See Also**
- [Example: Configuring BFD for OSPF on page 61](#)
  - *bfd-liveness-detection*

## Example: Configuring BFD for OSPF

This example shows how to configure the Bidirectional Forwarding Detection (BFD) protocol for OSPF.

- [Requirements on page 61](#)
- [Overview on page 62](#)
- [Configuration on page 63](#)
- [Verification on page 65](#)

### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See *Example: Configuring an OSPF Router Identifier*.
- Control OSPF designated router election. See *Example: Controlling OSPF Designated Router Election*.
- Configure a single-area OSPF network. See *Example: Configuring a Single-Area OSPF Network*.
- Configure a multiarea OSPF network. See *Example: Configuring a Multiarea OSPF Network*.
- Configure a multiarea OSPF network. See *Example: Configuring a Multiarea OSPF Network*.

---

## Overview

An alternative to adjusting the OSPF hello interval and dead interval settings to increase route convergence is to configure BFD. The BFD protocol is a simple hello mechanism that detects failures in a network. The BFD failure detection timers have shorter timer limits than the OSPF failure detection mechanisms, thereby providing faster detection.

BFD is useful on interfaces that are unable to detect failure quickly, such as Ethernet interfaces. Other interfaces, such as SONET interfaces, already have built-in failure detection. Configuring BFD on those interfaces is unnecessary.

You configure BFD on a pair of neighboring OSPF interfaces. Unlike the OSPF hello interval and dead interval settings, you do not have to enable BFD on all interfaces in an OSPF area.

In this example, you enable failure detection by including the **bfd-liveness-detection** statement on the neighbor OSPF interface **fe-0/1/0** in area 0.0.0.0 and configure the BFD packet exchange interval to 300 milliseconds, configure 4 as the number of missed hello packets that causes the originating interface to be declared down, and configure BFD sessions only for OSPF neighbors with full neighbor adjacency by including the following settings:

- **full-neighbors-only**—In Junos OS Release 9.5 and later, configures the BFD protocol to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors.
- **minimum-interval**—Configures the minimum interval, in milliseconds, after which the local routing device transmits hello packets as well as the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



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

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 no less than 500 ms. An interval of 1000 ms is recommended to avoid any instability issues.



**NOTE:**

- For the `bfdd` process, the detection time interval set is lower than 300 ms. If there is a high priority process such as `ppmd` running on the system, the CPU might spend time on the `ppmd` process rather than the `bfdd` process.
- For branch SRX Series devices, we recommend 1000 ms as the minimum keepalive time interval for BFD packets.

- 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 (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

- **multiplier**—Configures the number of hello packets not received by a neighbor that causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down. You can configure a value in the range from 1 through 255.

## Configuration

### CLI Quick Configuration

To quickly configure the BFD protocol for OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection minimum-interval
  300
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

**Step-by-Step Procedure** To configure the BFD protocol for OSPF on one neighboring interface:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
```

3. Specify the minimum transmit and receive intervals.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
```

4. Configure the number of missed hello packets that cause the originating interface to be declared down.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection multiplier 4
```

5. Configure BFD sessions only for OSPF neighbors with full neighbor adjacency.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```



**NOTE:** Repeat this entire configuration on the other neighboring interface.

**Results** Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0 {
    bfd-liveness-detection {
```



```

        minimum-interval 300;
        multiplier 4;
        full-neighbors-only;
    }
}

```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

---

Confirm that the configuration is working properly.

#### *Verifying the BFD Sessions*

**Purpose** Verify that the OSPF interfaces have active BFD sessions, and that session components have been configured correctly.

**Action** From operational mode, enter the **show bfd session detail** command.

**Meaning** The output displays information about the BFD sessions.

- The Address field displays the IP address of the neighbor.
- The Interface field displays the interface you configured for BFD.
- The State field displays the state of the neighbor and should show Full to reflect the full neighbor adjacency that you configured.
- The Transmit Interval field displays the time interval you configured to send BFD packets.
- The Multiplier field displays the multiplier you configured.

**See Also**

- [Understanding BFD for OSPF on page 59](#)
- *Understanding BFD Authentication for OSPF*
- *Example: Configuring BFD Authentication for OSPF*

### Configuring BFD for IS-IS

---

- [Understanding BFD for IS-IS on page 66](#)
- [Example: Configuring BFD for IS-IS on page 68](#)

## Understanding BFD for IS-IS

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than the failure detection mechanisms of IS-IS, providing faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (RX) interval by two if the local BFD instance is the reason for the session flap. The transmission (TX) interval is increased by two if the remote BFD instance is the reason for the session flap.

You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



**NOTE:** Starting with Junos OS Release 16.1R1, you can configure IS-IS BFD sessions for IPv6 by including the **bfd-liveness-detection** statement at the [edit protocols isis interface *interface-name* family inet[inet6] hierarchy level.

- For interfaces that support both IPv4 and IPv6 routing, the **bfd-liveness-detection** statement must be configured separately for each inet family.
- BFD over IPv6 link local address is currently not distributed because IS-IS uses link local addresses for forming adjacencies.
- BFD sessions over IPv6 must not have the same aggressive detection intervals as IPv4 sessions.
- BFD IPv6 sessions with detection intervals less than 2.5 seconds are currently not supported when nonstop active routing (NSR) is enabled.

To detect failures in the network, the set of statements in [Table 3 on page 66](#) are used in the configuration.

**Table 3: Configuring BFD for IS-IS**

Statement	Description
<b>bfd-liveness-detection</b>	Enable failure detection.

Table 3: Configuring BFD for IS-IS (continued)

Statement	Description
<b>minimum-interval milliseconds</b>	<p>Specify the minimum transmit and receive intervals for failure detection.</p> <p>This value represents the minimum interval at which the local router transmits hellos packets as well as the minimum interval at which the router expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.</p> <p><b>NOTE:</b> 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.</p> <p>Depending on your network environment, these additional recommendations might apply:</p> <ul style="list-style-type: none"> <li>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.</li> <li>For very large-scale network deployments with a large number of BFD sessions, please contact Juniper Networks customer support for more information.</li> <li>For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) 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.</li> </ul>
<b>minimum-receive-interval milliseconds</b>	<p>Specify only the minimum receive interval for failure detection.</p> <p>This value represents the minimum interval at which the local router expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number from 1 through 255,000 milliseconds.</p>
<b>multiplier number</b>	<p>Specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down.</p> <p>The default is 3, and you can configure a value from 1 through 225.</p>
<b>no-adaptation</b>	<p>Disable BFD adaptation.</p> <p>In Junos OS Release 9.0 and later, you can specify that the BFD sessions not adapt to changing network conditions.</p> <p><b>NOTE:</b> We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.</p>
<b>threshold</b>	<p>Specify the threshold for the following:</p> <ul style="list-style-type: none"> <li>Adaptation of the detection time           <p>When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.</p> </li> <li>Transmit interval</li> </ul> <p><b>NOTE:</b> The threshold value must be greater than the minimum transmit interval multiplied by the multiplier number.</p>

Table 3: Configuring BFD for IS-IS (continued)

Statement	Description
<code>transmit-interval</code> <code>minimum-interval</code>	Specify the minimum transmit interval for failure detection.  This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.
<code>version</code>	Specify the BFD version used for detection.  The default is to have the version detected automatically.



**NOTE:** You can trace BFD operations by including the `traceoptions` statement at the `[edit protocols bfd]` hierarchy level.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

- See Also**
- [Example: Configuring BFD for IS-IS on page 68](#)
  - [Understanding BFD Authentication for IS-IS](#)

## Example: Configuring BFD for IS-IS

This example describes how to configure the Bidirectional Forwarding Detection (BFD) protocol to detect failures in an IS-IS network.



**NOTE:** BFD is not supported with ISIS for IPV6 on QFX10000 series switches.

- [Requirements on page 68](#)
- [Overview on page 69](#)
- [Configuration on page 69](#)
- [Verification on page 72](#)

### Requirements

Before you begin, configure IS-IS on both routers. See *Example: Configuring IS-IS* for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

## Overview

This example shows two routers connected to each other. A loopback interface is configured on each router. IS-IS and BFD protocols are configured on both routers.

Figure 3 on page 69 shows the sample network.

Figure 3: Configuring BFD for IS-IS



## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Router R1

```

set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 5
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 2
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 3
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
  
```

#### Router R2

```

set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 6
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 3
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 4
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
  
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.



**NOTE:** To simply configure BFD for IS-IS, only the `minimum-interval` statement is required. The BFD protocol selects default parameters for all the other configuration statements when you use the `bfd-liveness-detection` statement without specifying any parameters.



**NOTE:** You can change parameters at any time without stopping or restarting the existing session. BFD automatically adjusts to the new parameter value. However, no changes to BFD parameters take place until the values resynchronize with each BFD peer.

To configure BFD for IS-IS on Routers R1 and R2:

1. Enable BFD failure detection for IS-IS.  

```
[edit protocols isis]  
user@R1# set interface so-0/0/0 bfd-liveness-detection  
  
[edit protocols isis]  
user@R2# set interface so-0/0/0 bfd-liveness-detection
```
2. Configure the threshold for the adaptation of the detection time, which must be greater than the multiplier number multiplied by the minimum interval.  

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R1# set detection-time threshold 5  
  
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R2# set detection-time threshold 6
```
3. Configure the minimum transmit and receive intervals for failure detection.  

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R1# set minimum-interval 2  
  
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R2# set minimum-interval 3
```
4. Configure only the minimum receive interval for failure detection.  

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R1# set minimum-receive-interval 1  
  
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R2# set minimum-receive-interval 1
```

5. Disable BFD adaptation.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set no-adaptation
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set no-adaptation
```

6. Configure the threshold for the transmit interval, which must be greater than the minimum transmit interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval threshold 3
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval threshold 4
```

7. Configure the minimum transmit interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval minimum-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval minimum-interval 1
```

8. Configure the multiplier number, which is the number of hello packets not received by the neighbor that causes the originating interface to be declared down.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set multiplier 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set multiplier 2
```

9. Configure the BFD version used for detection.

The default is to have the version detected automatically.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set version automatic
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set version automatic
```

## Results

From configuration mode, confirm your configuration by issuing the **show protocols isis interface** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols isis interface so-0/0/0
```

```
bfd-liveness-detection {
  version automatic;
  minimum-interval 2;
  minimum-receive-interval 1;
```

```
        multiplier 2;
        no-adaptation;
        transmit-interval {
            minimum-interval 1;
            threshold 3;
        }
        detection-time {
            threshold 5;
        }
    }
    ...
```

user@R2# show protocols isis interface so-0/0/0

```
    bfd-liveness-detection {
        version automatic;
        minimum-interval 3;
        minimum-receive-interval 1;
        multiplier 2;
        no-adaptation;
        transmit-interval {
            minimum-interval 1;
            threshold 4;
        }
        detection-time {
            threshold 6;
        }
    }
    ...
```

---

## Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1 and R2 on page 72](#)
- [Verifying That IS-IS Is Configured on page 73](#)
- [Verifying That BFD Is configured on page 74](#)

### ***Verifying the Connection Between Routers R1 and R2***

**Purpose** Make sure that Routers R1 and R2 are connected to each other.

**Action** Ping the other router to check the connectivity between the two routers as per the network topology.

user@R1> ping 10.0.0.2

```
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.367 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.662 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=1.291 ms
^C
--- 10.0.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.291/1.440/1.662/0.160 ms
```



```

user@R2> ping 10.0.0.1

PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.287 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.310 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.289 ms
^C
--- 10.0.0.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.287/1.295/1.310/0.010 ms

```

**Meaning** Routers R1 and R2 are connected to each other.

### *Verifying That IS-IS Is Configured*

**Purpose** Make sure that the IS-IS instance is running on both routers.

**Action** Use the **show isis database** statement to check if the IS-IS instance is running on both routers, R1 and R2.

```

user@R1> show isis database

IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4a571  0x30c5    1195 L1 L2
R2.00-00        0x4a586  0x4b7e    1195 L1 L2
R2.02-00        0x330ca1 0x3492    1196 L1 L2
  3 LSPs

IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4a856  0x5db0    1194 L1 L2
R2.00-00        0x4a89d  0x149b    1194 L1 L2
R2.02-00        0x1fb2ff 0xd302    1194 L1 L2
  3 LSPs

user@R2> show isis database

IS-IS level 1 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4b707  0xcc80    1195 L1 L2
R2.00-00        0x4b71b  0xeb37    1198 L1 L2
R2.02-00        0x33c2ce 0xb52d    1198 L1 L2
  3 LSPs

IS-IS level 2 link-state database:
LSP ID          Sequence Checksum Lifetime Attributes
R1.00-00        0x4b9f2  0xee70    1192 L1 L2
R2.00-00        0x4ba41  0x9862    1197 L1 L2
R2.02-00        0x3      0x6242    1198 L1 L2
  3 LSPs

```

**Meaning** IS-IS is configured on both routers, R1 and R2.

### *Verifying That BFD Is configured*

**Purpose** Make sure that the BFD instance is running on both routers, R1 and R2.

**Action** Use the **show bfd session detail** statement to check if BFD instance is running on the routers.

user@R1> **show bfd session detail**

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	so-0/0/0	2.000	1.000	2
Client ISIS R2, TX interval 0.001, RX interval 0.001					
Client ISIS R1, TX interval 0.001, RX interval 0.001					
Session down time 00:00:00, previous up time 00:00:15					
Local diagnostic NbrSignal, remote diagnostic NbrSignal					
Remote state AdminDown, version 1					
Router 3, routing table index 17					

1 sessions, 2 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

user@R2> **show bfd session detail**

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.1	Up	so-0/0/0	2.000	1.000	2
Client ISIS R2, TX interval 0.001, RX interval 0.001					
Session down time 00:00:00, previous up time 00:00:05					
Local diagnostic NbrSignal, remote diagnostic NbrSignal					
Remote state AdminDown, version 1					
Router 2, routing table index 15					

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

**Meaning** BFD is configured on Routers R1 and R2 for detecting failures in the IS-IS network.

**See Also**

- [Understanding BFD for IS-IS on page 66](#)

---

## Configuring BFD for RIP

- [Understanding BFD for RIP on page 75](#)
- [Example: Configuring BFD for RIP on page 75](#)

## Understanding BFD for RIP

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. BFD failure detection times are shorter than RIP detection times, providing faster reaction times to various kinds of failures in the network. Instead of waiting for the routing protocol neighbor timeout, BFD provides rapid detection of link failures. BFD timers are adaptive and can be adjusted to be more or less aggressive. For example, a timer can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured. Note that the functionality of configuring BFD for RIP described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

BFD enables quick failover between a primary and a secondary routed path. The protocol tests the operational status of the interface multiple times per second. BFD provides for configuration timers and thresholds for failure detection. For example, if the minimum interval is set for 50 milliseconds and the threshold uses the default value of three missed messages, a failure is detected on an interface within 200 milliseconds of the failure.

Intervening devices (for example, an Ethernet LAN switch) hide link-layer failures from routing protocol peers, such as when two routers are connected by way of a LAN switch, where the local interface status remains up even when a physical fault happens on the remote link. Link-layer failure detection times vary, depending on the physical media and the Layer 2 encapsulation. BFD can provide fast failure detection times for all media types, encapsulations, topologies, and routing protocols.

To enable BFD for RIP, both sides of the connection must receive an update message from the peer. By default, RIP does not export any routes. Therefore, you must enable update messages to be sent by configuring an export policy for routes before a BFD session is triggered.

**See Also** • [Example: Configuring BFD for RIP on page 75](#)

## Example: Configuring BFD for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) for a RIP network.

- [Requirements on page 75](#)
- [Overview on page 76](#)
- [Configuration on page 78](#)
- [Verification on page 80](#)

---

### Requirements

No special configuration beyond device initialization is required before configuring this example.

## Overview

---

To enable failure detection, include the **bfd-liveness-detection** statement:

```
bfd-liveness-detection {  
  detection-time {  
    threshold milliseconds;  
  }  
  minimum-interval milliseconds;  
  minimum-receive-interval milliseconds;  
  multiplier number;  
  no-adaptation;  
  transmit-interval {  
    threshold milliseconds;  
    minimum-interval milliseconds;  
  }  
  version (1 | automatic);  
}
```

Optionally, you can specify the threshold for the adaptation of the detection time by including the **threshold** statement. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.

To specify the minimum transmit and receive interval for failure detection, include the **minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval at which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. This examples sets a minimum interval of 600 milliseconds.



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

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 (NSR) 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.
-

You can optionally specify the minimum transmit and receive intervals separately.

To specify only the minimum receive interval for failure detection, include the **minimum-receive-interval** statement. This value represents the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255.00 milliseconds.

To specify only the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify the number of hello packets not received by a neighbor that causes the originating interface to be declared down, include the **multiplier** statement. The default is 3, and you can configure a value in the range from 1 through 255.

To specify the threshold for detecting the adaptation of the transmit interval, include the **transmit-interval threshold** statement. The threshold value must be greater than the transmit interval.


To specify the BFD version used for detection, include the **version** statement. The default is to have the version detected automatically.

You can trace BFD operations by including the **traceoptions** statement at the **[edit protocols bfd]** hierarchy level.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions. To disable BFD adaptation, include the **no-adaptation** statement. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

Figure 4 on page 77 shows the topology used in this example.

Figure 4: RIP BFD Network Topology



“CLI Quick Configuration” on page 78 shows the configuration for all of the devices in Figure 4 on page 77. The section “Step-by-Step Procedure” on page 78 describes the steps on Device R1.

## Configuration

---

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R2**

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD for a RIP network:

1. Configure the network interfaces.

```
[edit interfaces]
```

```
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Enable BFD.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600
```

6. Configure tracing operations to track BFD messages.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}

user@R1# show protocols
bfd {
  traceoptions {
    file bfd-trace;
    flag all;
  }
}
```

```
rip {
  group rip-group {
    export advertise-routes-through-rip;
    bfd-liveness-detection {
      minimum-interval 600;
    }
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

---

### Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Up on page 80](#)
- [Checking the BFD Trace File on page 80](#)

#### *Verifying That the BFD Sessions Are Up*

**Purpose** Make sure that the BFD sessions are operating.

**Action** From operational mode, enter the **show bfd session** command.

```
user@R1> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

```
1 sessions, 1 clients
Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps
```

**Meaning** The output shows that there are no authentication failures.

#### *Checking the BFD Trace File*

**Purpose** Use tracing operations to verify that BFD packets are being exchanged.

**Action** From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
```



```

Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 6f
...

```

**Meaning** The output shows the normal functioning of BFD.

**See Also** • [Understanding BFD for RIP on page 75](#)

## Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

This example shows the global configuration.

```

[edit protocols pim]
user@host# edit interface fe-1/0/0.0 family inet bfd-liveness-detection

```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-receive-interval 350
```

4. (Optional) Configure other BFD settings.

As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-interval 350
```

5. Configure the threshold for the adaptation of the BFD session detection time.

When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set detection-time threshold 800
```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set multiplier 50
```

7. Configure the BFD version.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set version 1
```

8. Specify that BFD sessions should not adapt to changing network conditions.

We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set no-adaptation
```

9. Verify the configuration by checking the output of the **show bfd session** command.

Related Documentation • [show bfd session](#)

## Configuring Independent Micro BFD Sessions for LAG

- [Understanding Independent Micro BFD Sessions for LAG on page 83](#)
- [Configuring Independent Micro BFD Sessions for LAG on page 85](#)
- [Example: Configuring Independent Micro BFD Sessions for LAG on page 91](#)

### Understanding Independent Micro BFD Sessions for LAG

Starting with Junos OS Release 13.3, this feature is supported on the following PIC/FPC types:

- PC-1XGE-XENPAK (Type 3 FPC)
- PD-4XGE-XFP (Type 4 FPC)
- PD-5-10XGE-SFPP (Type 4 FPC)
- 24x10GE (LAN/WAN) SFPP, 12x10GE (LAN/WAN) SFPP, 1x100GE Type 5 PICs
- All MPCs on MX Series with Ethernet MICs
- FPC-PTX-P1-A on PTX5000 with 10-Gigabit Ethernet interfaces
- FPC2-PTX-P1A on PTX5000 with 10-Gigabit Ethernet interfaces in Junos OS Release 14.1 and later
- All FPCs on PTX Series with Ethernet interfaces in Junos OS Release 14.1R3 and later 14.1 releases, and Junos 14.2 and later



**TIP:** See *PTX Series PIC/FPC Compatibility* for a list of PICs that are supported on each PTX Series FPC.

The Bidirectional Forwarding Detection (BFD) protocol is a simple detection protocol that quickly detects failures in the forwarding paths. A link aggregation group (LAG) combines multiple links between devices that are in point-to-point connections, thereby increasing bandwidth, providing reliability, and allowing load balancing. To run a BFD session on LAG interfaces, configure an independent, asynchronous mode BFD session on every LAG member link in a LAG bundle. Instead of a single BFD session monitoring the status of the UDP port, independent micro BFD sessions monitor the status of individual member links.

The individual BFD sessions determine the Layer 2 and Layer 3 connectivity of each member link in the LAG. Once a BFD session is established on a particular link, the member links are attached to the LAG and the load balancer either by a static configuration or by the Link Aggregation Control Protocol (LACP). If the member links are attached to the LAG by a static configuration, the device control process acts as the client to the micro BFD session. When member links are attached to the LAG by the LACP, the LACP acts as the client to the micro BFD session.

When the micro BFD session is up, a LAG link is established and data is transmitted over that LAG link. If the micro BFD session on a member link is down, that particular member link is removed from the load balancer, and the LAG managers stop directing traffic to that link. These micro BFD sessions are independent of each other despite having a single client that manages the LAG interface.



**NOTE:** Starting with Junos OS Release 13.3, IANA has allocated 01-00-5E-90-00-01 as the dedicated MAC address for micro BFD. Dedicated MAC mode is used by default for micro BFD sessions, in accordance with the latest draft for BFD over LAG.

Micro BFD sessions run in the following modes:

- Distribution Mode—Micro BFD sessions are distributed by default at Layer 3.
- Non-Distribution Mode—You can configure the BFD session to run in this mode by including the **no-delegate-processing** statement under periodic packet management (PPM). In this mode, the packets are being sent or received by the Routing Engine at Layer 2.

A pair of routing devices in a LAG exchange BFD packets at a specified, regular interval. The routing device detects a neighbor failure when it stops receiving a reply after a specified interval. This allows the quick verification of member link connectivity with or without LACP. A UDP port distinguishes BFD over LAG packets from BFD over single-hop IP.



**NOTE:** IANA has allocated 6784 as the UDP destination port for micro BFD.

To enable failure detection for LAG networks for aggregated Ethernet interfaces:

- Include the **bfd-liveness-detection** statement in the configuration.
- Specify a hold-down interval value to set the minimum time that the BFD session must remain up before a state change notification is sent to the other members in the LAG network.
- Specify the minimum interval that indicates the time interval for transmitting and receiving data.
- Starting with Junos OS Release 14.1, specify the neighbor in a BFD session. In releases prior to Junos OS Release 16.1, you must configure the loopback address of the remote destination as the neighbor address. Beginning with Junos OS Release 16.1, you can also configure this feature with aggregated Ethernet interface address of the remote destination as the neighbor address.



**CAUTION:** Deactivate **bfd-liveness-detection** at the [edit interfaces aex aggregated-ether-options] hierarchy level or deactivate the aggregated Ethernet interface before changing the neighbor address from loopback

IP address to aggregated Ethernet interface IP address. Modifying the local and neighbor address without deactivating `bfd-liveness-detection` or the aggregated Ethernet interface first might cause micro BFD sessions failure.



**NOTE:** Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD local-address against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.



**NOTE:** This feature works only when both the devices support BFD. If BFD is configured at one end of the LAG, this feature does not work.

For the IPv6 address family, disable duplicate address detection before configuring this feature with AE interface addresses. To disable duplicate address detection, include the `dad-disable` statement at the `[edit interface aex unit y family inet6]` hierarchy level.

- See Also
- [authentication on page 108](#)
  - [bfd-liveness-detection on page 109](#)
  - [detection-time on page 111](#)
  - [transmit-interval on page 114](#)
  - [Configuring Independent Micro BFD Sessions for LAG on page 85](#)
  - [Example: Configuring Independent Micro BFD Sessions for LAG on page 91](#)

## Configuring Independent Micro BFD Sessions for LAG

The Bidirectional Forwarding Detection (BFD) protocol is a simple detection protocol that quickly detects failures in the forwarding paths. A link aggregation group (LAG) combines multiple links between devices that are in point-to-point connections, thereby increasing bandwidth, providing reliability, and allowing load balancing. To run a BFD session on LAG interfaces, configure an independent, asynchronous mode BFD session on every LAG member link in a LAG bundle. Instead of a single BFD session monitoring the status of the UDP port, independent micro BFD sessions monitor the status of individual member links.

To enable failure detection for aggregated Ethernet interfaces:

1. Include the following statement in the configuration at the `[edit interfaces aex aggregated-ether-options]` hierarchy level:

```
bfd-liveness-detection {
```

```
authentication {  
    algorithm algorithm-name;  
    key-chain key-chain-name;  
    loose-check;  
}  
detection-time {  
    threshold milliseconds;  
}  
holddown-interval milliseconds;  
local-address bfd-local-address;  
minimum-interval milliseconds;  
minimum-receive-interval milliseconds;  
multiplier number;  
neighbor bfd-neighbor-address;  
no-adaptation;  
transmit-interval {  
    minimum-interval milliseconds;  
    threshold milliseconds;  
}  
version (1 | automatic);  
}
```

2. Configure the authentication criteria of the BFD session for LAG.

To specify the authentication criteria, include the **authentication** statement:

```
bfd-liveness-detection {  
    authentication {  
        algorithm algorithm-name;  
        key-chain key-chain-name;  
        loose-check;  
    }  
}
```

- Specify the algorithm to be used to authenticate the BFD session. You can use one of the following algorithms for authentication:
    - keyed-md5
    - keyed-sha-1
    - meticulous-keyed-md5
    - meticulous-keyed-sha-1
    - simple-password
  - To configure the key chain, specify the name that is associated with the security key for the BFD session. The name you specify must match one of the key chains configured in the **authentication-key-chains** *key-chain* statement at the **[edit security]** hierarchy level.
  - Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.
3. Configure BFD timers for aggregated Ethernet interfaces.

To specify the BFD timers, include the **detection-time** statement:

```
bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
}
```

Specify the threshold value. This is the maximum time interval for detecting a BFD neighbor. If the transmit interval is greater than this value, the device triggers a trap.

4. Configure a hold-down interval value to set the minimum time that the BFD session must remain up before a state change notification is sent to the other members in the LAG network.

To specify the hold-down interval, include the **holddown-interval** statement:

```
bfd-liveness-detection {
  holddown-interval milliseconds;
}
```

You can configure a number in the range from 0 through 255,000 milliseconds, and the default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

This value represents the minimum interval at which the local routing device transmits BFD packets, as well as the minimum interval in which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.

5. Configure the source address for the BFD session.

To specify a local address, include the **local-address** statement:

```
bfd-liveness-detection {
  local-address bfd-local-address;
}
```

The BFD local address is the loopback address of the source of the BFD session.



**NOTE:** Beginning with Junos OS Release 16.1, you can also configure this feature with the AE interface address as the local address in a micro BFD session. For the IPv6 address family, disable duplicate address detection before configuring this feature with the AE interface address. To disable duplicate address detection, include the **dad-disable** statement at the **[edit interface aex unit y family inet6]** hierarchy level.

Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD **local-address** against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

6. Specify the minimum interval that indicates the time interval for transmitting and receiving data.

This value represents the minimum interval at which the local routing device transmits BFD packets, as well as the minimum interval in which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement:

```
bfd-liveness-detection {  
    minimum-interval milliseconds;  
}
```



**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.

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.

7. Specify only the minimum receive interval for failure detection by including the **minimum-receive-interval** statement:

```
bfd-liveness-detection {  
    minimum-receive-interval milliseconds;  
}
```

This value represents the minimum interval in which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds.

8. Specify the number of BFD packets that were not received by the neighbor that causes the originating interface to be declared down by including the **multiplier** statement:



```
bfd-liveness-detection {
  multiplier number;
}
```

The default value is 3. You can configure a number in the range from 1 through 255.

9. Configure the neighbor in a BFD session.

The neighbor address can be either an IPv4 or an IPv6 address.

To specify the next hop of the BFD session, include the **neighbor** statement:

```
bfd-liveness-detection {
  neighbor bfd-neighbor-address;
}
```

The BFD neighbor address is the loopback address of the remote destination of the BFD session.



**NOTE:** Beginning with Junos OS Release 16.1, you can also configure the AE interface address of the remote destination as the BFD neighbor address in a micro BFD session.

10. (Optional) Configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement:

```
bfd-liveness-detection {
  no-adaptation;
}
```



**NOTE:** We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

11. Specify a threshold for detecting the adaptation of the detection time by including the **threshold** statement:

```
bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
}
```

When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the minimum-interval or the minimum-receive-interval value. The threshold must be a higher value than the multiplier for either of these configured values. For example, if the minimum-receive-interval is 300 ms and the multiplier is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value greater than 900.

12. Specify only the minimum transmit interval for failure detection by including the **transmit-interval minimum-interval** statement:

```
bfd-liveness-detection {  
  transmit-interval {  
    minimum-interval milliseconds;  
  }  
}
```

This value represents the minimum interval at which the local routing device transmits BFD packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

13. Specify the transmit threshold for detecting the adaptation of the transmit interval by including the **transmit-interval threshold** statement:

```
bfd-liveness-detection {  
  transmit-interval {  
    threshold milliseconds;  
  }  
}
```

The threshold value must be greater than the transmit interval. When the BFD session detection time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the minimum-interval or the minimum-receive-interval value. The threshold must be a higher value than the multiplier for either of these configured values.

14. Specify the BFD version by including the **version** statement:

```
bfd-liveness-detection {  
  version (1 | automatic);  
}
```

The default is to have the version detected automatically.



**NOTE:** This feature works when both the devices support BFD. If BFD is configured at only one end of the LAG, this feature does not work.

---

- See Also**
- [authentication on page 108](#)
  - [bfd-liveness-detection on page 109](#)
  - [detection-time on page 111](#)
  - [Example: Configuring Independent Micro BFD Sessions for LAG on page 91](#)
  - [Understanding Independent Micro BFD Sessions for LAG on page 83](#)

## Example: Configuring Independent Micro BFD Sessions for LAG

This example shows how to configure an independent micro BFD session for aggregated Ethernet interfaces.

- [Requirements on page 91](#)
- [Overview on page 91](#)
- [Configuration on page 92](#)
- [Verification on page 97](#)

### Requirements

This example uses the following hardware and software components:

- MX Series routers with Junos Trio chipset
- T Series routers with Type 4 FPC or Type 5 FPC

BFD for LAG is supported on the following PIC types on T-Series:

- PC-1XGE-XENPAK (Type 3 FPC),
- PD-4XGE-XFP (Type 4 FPC),
- PD-5-10XGE-SFPP (Type 4 FPC),
- 24x10GE (LAN/WAN) SFPP, 12x10GE (LAN/WAN) SFPP, 1X100GE Type 5 PICs
- PTX Series routers with 24X10GE (LAN/WAN) SFPP
- Junos OS Release 13.3 or later running on all devices

### Overview

The example includes two routers that are directly connected. Configure two aggregated Ethernet interfaces, AE0 for IPv4 connectivity and AE1 for IPv6 connectivity. Configure micro BFD session on the AE0 bundle using IPv4 addresses as local and neighbor endpoints on both routers. Configure micro BFD session on the AE1 bundle using IPv6 addresses as local and neighbor endpoints on both routers. This example verifies that independent micro BFD sessions are active in the output.

### Topology

[Figure 5 on page 91](#) shows the sample topology.

**Figure 5: Configuring an Independent Micro BFD Session for LAG**



## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Router R0**

```

set interfaces ge-1/0/1 unit 0 family inet address 20.20.20.1/30
set interfaces ge-1/0/1 unit 0 family inet6 address 3ffe::1:1/126
set interfaces xe-4/0/0 gigether-options 802.3ad ae0
set interfaces xe-4/0/1 gigether-options 802.3ad ae0
set interfaces xe-4/1/0 gigether-options 802.3ad ae1
set interfaces xe-4/1/1 gigether-options 802.3ad ae1
set interfaces lo0 unit 0 family inet address 10.255.106.107/32
set interfaces lo0 unit 0 family inet6 address 201:DB8:251::aa:aa:1/126
set interfaces ae0 aggregated-ether-options bfd-liveness-detection minimum-interval
  100
set interfaces ae0 aggregated-ether-options bfd-liveness-detection neighbor
  10.255.106.102
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.107
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 10.0.0.1/30
set interfaces ae1 aggregated-ether-options bfd-liveness-detection minimum-interval
  100
set interfaces ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae1 aggregated-ether-options bfd-liveness-detection neighbor
  201:DB8:251::bb:bb:1
set interfaces ae1 aggregated-ether-options bfd-liveness-detection local-address
  201:DB8:251::aa:aa:1
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet6 address 5555::1/126
set interface ae1 unit 0 family inet6 dad-disable
set routing-options nonstop-routing
set routing-options static route 30.30.30.0/30 next-hop 10.0.0.2
set routing-options rib inet6.0 static route 3ffe::1:2/126 next-hop 5555::2
set protocols bfd traceoptions file bfd
set protocols bfd traceoptions file size 100m
set protocols bfd traceoptions file files 10
set protocols bfd traceoptions flag all

```

**Router R1**

```

set interfaces ge-1/1/8 unit 0 family inet address 30.30.30.1/30
set interfaces ge-1/1/8 unit 0 family inet6 address 3ffe::1:2/126
set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set interfaces xe-0/0/2 gigether-options 802.3ad ae1
set interfaces xe-0/0/3 gigether-options 802.3ad ae1
set interfaces lo0 unit 0 family inet address 10.255.106.102/32
set interfaces lo0 unit 0 family inet6 address 201:DB8:251::bb:bb:1/126

```

```

set interfaces ae0 aggregated-ether-options bfd-liveness-detection minimum-interval
  150
set interfaces ae0 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae0 aggregated-ether-options bfd-liveness-detection neighbor
  10.255.106.107
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.102
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 family inet address 10.0.0.2/30
set interfaces ae1 aggregated-ether-options bfd-liveness-detection minimum-interval
  200
set interfaces ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae1 aggregated-ether-options bfd-liveness-detection neighbor
  201:DB8:251::aa:aa:1
set interfaces ae1 aggregated-ether-options bfd-liveness-detection local-address
  201:DB8:251::bb:bb:1
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp passive
set interfaces ae1 unit 0 family inet6 address 5555::2/126
set routing-options static route 20.20.20.0/30 next-hop 10.0.0.1
set routing-options rib inet6.0 static route 3ffe::1/126 next-hop 5555::1

```

### Configuring a Micro BFD Session for Aggregated Ethernet Interfaces

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *“Using the CLI Editor in Configuration Mode”* in the *CLI User Guide*.



**NOTE:** Repeat this procedure for Router R1, modifying the appropriate interface names, addresses, and any other parameters for each router.

To configure a micro BFD session for aggregated Ethernet interfaces on Router R0:

1. Configure the physical interfaces.

```

[edit interfaces]
user@R0# set ge-1/0/1 unit 0 family inet address 20.20.20.1/30
user@R0# set ge-1/0/1 unit 0 family inet6 address 3ffe::1/126
user@R0# set xe-4/0/0 gigether-options 802.3ad ae0
user@R0# set xe-4/0/1 gigether-options 802.3ad ae0
user@R0# set xe-4/1/0 gigether-options 802.3ad ae1
user@R0# set xe-4/1/1 gigether-options 802.3ad ae1

```

2. Configure the loopback interface.

```

[edit interfaces]
user@R0# set lo0 unit 0 family inet address 10.255.106.107/32
user@R0# set lo0 unit 0 family inet6 address 201:DB8:251::aa:aa:1/128

```

3. Configure an IP address on the aggregated Ethernet interface ae0 with either IPv4 or IPv6 addresses, as per your network requirements.

[edit interfaces]

```
user@R0# set ae0 unit 0 family inet address 10.0.0.1/30
```

4. Set the routing option, create a static route, and set the next-hop address.



**NOTE:** You can configure either an IPv4 or IPv6 static route, depending on your network requirements.

[edit routing-options]

```
user@R0# set nonstop-routing
```

```
user@R0# set static route 30.30.30.0/30 next-hop 10.0.0.2
```

```
user@R0# set rib inet6.0 static route 3ffe::1:2/126 next-hop 5555::2
```

5. Configure the Link Aggregation Control Protocol (LACP).

[edit interfaces]

```
user@R0# set ae0 aggregated-ether-options lacp active
```

6. Configure BFD for the aggregated Ethernet interface ae0, and specify the minimum interval, local IP address, and the neighbor IP address.

[edit interfaces]

```
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection
  minimum-interval 100
```

```
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection multiplier 3
```

```
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection neighbor
  10.255.106.102
```

```
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.107
```

```
user@R0# set ae0 aggregated-ether-options minimum-links 1
```

```
user@R0# set ae0 aggregated-ether-options link-speed 10g
```

7. Configure an IP address on the aggregated Ethernet interface ae1.

You can assign either IPv4 or IPv6 addresses as per your network requirements.

[edit interfaces]

```
user@R0# set ae1 unit 0 family inet6 address 5555::1/126
```

8. Configure BFD for the aggregated Ethernet interface ae1.

[edit interfaces]

```
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection
  minimum-interval 100
```

```
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
```

```
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection neighbor
  201:DB8:251::bb:bb:1
```

```
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection local-address
  201:DB8:251::aa:aa:1
```

```
user@R0# set ae1 aggregated-ether-options minimum-links 1
user@R0# set ae1 aggregated-ether-options link-speed 10g
```



**NOTE:** Beginning with Junos OS Release 16.1, you can also configure this feature with the AE interface address as the local address in a micro BFD session.

Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD local-address against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

9. Configure tracing options for BFD for troubleshooting.

```
[edit protocols]
user@R0# set bfd traceoptions file bfd
user@R0# set bfd traceoptions file size 100m
user@R0# set bfd traceoptions file files 10
user@R0# set bfd traceoptions flag all
```

## Results

From configuration mode, enter the **show interfaces**, **show protocols**, and **show routing-options** commands and confirm your configuration. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0> show interfaces
traceoptions {
  flag bfd-events;
}
ge-1/0/1 {
  unit 0 {
    family inet {
      address 20.20.20.1/30;
    }
    family inet6 {
      address 3ffe::1/126;
    }
  }
}
xe-4/0/0 {
  enable;
  gigether-options {
    802.3ad ae0;
  }
}
xe-4/0/1 {
  gigether-options {
    802.3ad ae0;
```

```
    }
  }
  xe-4/1/0 {
    enable;
    gigether-options {
      802.3ad ae1;
    }
  }
  xe-4/1/1 {
    gigether-options {
      802.3ad ae1;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.106.107/32;
      }
      family inet6 {
        address 201:DB8:251::aa:aa:1/128;
      }
    }
  }
  ae0 {
    aggregated-ether-options {
      bfd-liveness-detection {
        minimum-interval 100;
        neighbor 10.255.106.102;
        local-address 10.255.106.107;
      }
      minimum-links 1;
      link-speed 10g;
      lacp {
        active;
      }
    }
    unit 0 {
      family inet {
        address 10.0.0.1/30;
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      bfd-liveness-detection {
        minimum-interval 100;
        multiplier 3;
        neighbor 201:DB8:251::bb:bb:1;
        local-address 201:DB8:251::aa:aa:1;
      }
      minimum-links 1;
      link-speed 10g;
    }
    unit 0 {
      family inet6 {
        address 5555::1/126;
      }
    }
  }
}
```



```

    }
  }
}

user@R0> show protocols
bfd {
  traceoptions {
    file bfd size 100m files 10;
    flag all;
  }
}

user@R0> show routing-options
nonstop-routing ;
rib inet6.0 {
  static {
    route 3ffe:1:2/126 {
      next-hop 5555::2;
    }
  }
}
static {
  route 30.30.30.0/30 {
    next-hop 10.0.0.2;
  }
}
}

```

If you are done configuring the device, commit the configuration.

```
user@R0# commit
```

### Verification

Confirm that the configuration is working properly.

- [Verifying That the Independent BFD Sessions Are Up on page 97](#)
- [Viewing Detailed BFD Events on page 99](#)

#### *Verifying That the Independent BFD Sessions Are Up*

**Purpose** Verify that the micro BFD sessions are up, and view details about the BFD sessions.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@R0> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.255.106.102	Up	xe-4/0/0	9.000	3.000	3

```
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13, previous down time 00:00:06
Local diagnostic None, remote diagnostic None
Remote heard, hears us, version 1
Replicated
Session type: Micro BFD
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 0.100, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 21, remote discriminator 75
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x0
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.255.106.102	Up	xe-4/0/1	9.000	3.000	3

```
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13, previous down time 00:00:07
Local diagnostic None, remote diagnostic None
Remote heard, hears us, version 1
Replicated
Session type: Micro BFD
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 0.100, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 19, remote discriminator 74
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x0
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
201:DB8:251::bb:bb:1	Up	xe-4/1/1	9.000	3.000	3

```
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13
Local diagnostic None, remote diagnostic None
Remote not heard, hears us, version 1
Replicated
Session type: Micro BFD
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 17, remote discriminator 67
Echo mode disabled/inactive, no-absorb, no-refresh
Remote is control-plane independent
Session ID: 0x0
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
201:DB8:251::bb:bb:1 3		UP	xe-4/1/0	9.000	3.000

Client LACPD, TX interval 0.100, RX interval 0.100  
 Session up time 4d 23:13  
 Local diagnostic None, remote diagnostic None  
 Remote not heard, hears us, version 1  
 Replicated  
 Session type: **Micro BFD**  
 Min async interval 0.100, min slow interval 1.000  
 Adaptive async TX interval 0.100, RX interval 0.100  
 Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3  
 Remote min TX interval 3.000, min RX interval 3.000, multiplier 3  
 Local discriminator 16, remote discriminator 66  
 Echo mode disabled/inactive, no-absorb, no-refresh  
 Remote is control-plane independent  
 Session ID: 0x0

4 sessions, 4 clients  
 Cumulative transmit rate 2.0 pps, cumulative receive rate 1.7 pps

**Meaning** The Micro BFD field represents the independent micro BFD sessions running on the links in a LAG. The TX interval *item*, RX interval *item* output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under **bfd-liveness-detection** statement.

#### Viewing Detailed BFD Events

**Purpose** View the contents of the BFD trace file to assist in troubleshooting, if required.

**Action** From operational mode, enter the **file show /var/log/bfd** command.

```

user@R0> file show /var/log/bfd
Jun  5 00:48:59   Protocol (1) len 1: BFD
Jun  5 00:48:59   Data (9) len 41: (hex) 42 46 44 20 6e 65 69 67 68 62 6f 72 20
31 30 2e 30 2e 30
Jun  5 00:48:59 PPM Trace: BFD neighbor 10.255.106.102 (IFL 349) set, 9 0
Jun  5 00:48:59 Received Downstream RcvPkt (19) len 108:
Jun  5 00:48:59   IfIndex (3) len 4: 329
Jun  5 00:48:59   Protocol (1) len 1: BFD
Jun  5 00:48:59   SrcAddr (5) len 8: 10.255.106.102
Jun  5 00:48:59   Data (9) len 24: (hex) 00 88 03 18 00 00 00 4b 00 00 00 15 00
2d c6 c0 00 2d c6
Jun  5 00:48:59   PktError (26) len 4: 0
Jun  5 00:48:59   RtblIdx (24) len 4: 0
Jun  5 00:48:59   MultiHop (64) len 1: (hex) 00
Jun  5 00:48:59   Unknown (168) len 1: (hex) 01
Jun  5 00:48:59   Unknown (171) len 2: (hex) 02 3d
Jun  5 00:48:59   Unknown (172) len 6: (hex) 80 71 1f c7 81 c0
Jun  5 00:48:59   Authenticated (121) len 1: (hex) 01
Jun  5 00:48:59 BFD packet from 10.0.0.2 (IFL 329), len 24
Jun  5 00:48:59   Ver 0, diag 0, mult 3, len 24
Jun  5 00:48:59   Flags: IHU Fate
Jun  5 00:48:59   My discr 0x0000004b, your discr 0x00000015

```

```
Jun  5 00:48:59    Tx ivl 3000000, rx ivl 3000000, echo rx ivl 0
Jun  5 00:48:59 [THROTTLE]bfd_rate_limit_can_accept_pkt: session 10.255.106.102
is up or already in program thread
Jun  5 00:48:59 Replicate: marked session (discr 21) for update
```

**Meaning** BFD messages are being written to the specified trace file.

- See Also**
- [authentication on page 108](#)
  - [bfd-liveness-detection on page 109](#)
  - [detection-time on page 111](#)
  - [Configuring Independent Micro BFD Sessions for LAG on page 85](#)
  - [Understanding Independent Micro BFD Sessions for LAG on page 83](#)

---

## Understanding Distributed BFD

Bidirectional Forwarding Detection (BFD) is a protocol to verify the liveness of data path.

The terms *nondistributed BFD* and *centralized BFD* refer to BFD that runs on the Routing Engine. The term *distributed BFD* refers to BFD that runs on the Packet Forwarding Engine.



**NOTE:** By default, SRX Series devices operate in centralized BFD mode.

- **Single-hop BFD**—Single-hop BFD in Junos OS runs in distributed mode by default. The exceptions are OSPFv3 BFD and PIMv6 BFD, for which only nondistributed BFD is supported. Single-hop BFD control packets use UDP port 3784.
- **Multihop BFD**—One desirable application of BFD is to detect connectivity to routing devices that span multiple network hops and follow unpredictable paths. This is known as a multihop session. Prior to Junos OS Release 12.3, multihop BFD is nondistributed and runs on the Routing Engine. Starting in Junos OS Release 12.3, multihop BFD runs in distributed mode by default. Multihop BFD control packets use UDP port 4784.



**NOTE:** In a multichassis link aggregation group setup, Inter-Chassis Control Protocol (ICCP) uses BFD in multihop mode. Multihop BFD runs in centralized mode in this kind of setup prior to Junos OS Release 12.3 and continues to do so as of Junos OS Release 12.3 and later.

For both single-hop BFD and multihop BFD, the BFD session can be made to run on the Routing Engine (in nondistributed mode) by configuring **set routing-options ppm no-delegate-processing** and then running the **clear bfd session** command.

The benefits of distributed BFD are mainly in the scaling and performance areas.

The benefits are as follows:

- Allows for the creation of a larger number of BFD sessions.
- Runs BFD sessions with a shorter transfer/receive timer interval, which can in turn be used to bring down the overall detection time.
- Separates the functionality of BFD from that of the Routing Engine. This means that a BFD session can stay up during graceful restart, even with an aggressive interval. The minimum interval for Routing Engine-based BFD sessions to survive graceful Routing Engine switchover is 2500 ms. This is improved to sub-second times with distribution.
- Offloads the processing to the FPC CPU. This frees up the Routing Engine CPU, resulting in improved scaling and performance for Routing Engine-based applications.
- Starting with Junos OS Release 15.1X49-D100, dedicated BFD is supported on SRX340, SRX345, and SRX1500 devices.

Starting with Junos OS Release 15.1X49-D100, real-time BFD is supported on SRX300 and SRX320 devices.

Starting with Junos OS Release 15.1X49-D110, dedicated BFD is supported on SRX550M devices.

Starting with Junos OS Release 12.3X48-D60, dedicated BFD is supported on SRX240, SRX550, and SRX650 devices.

Starting with Junos OS Release 12.3X48-D60, real-time BFD is supported on SRX100, SRX110, SRX210, and SRX220 devices.

- To enable dedicated BFD on SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, SRX650, and SRX1500 devices, use the **set chassis dedicated-ukern-cpu** command.

Enabling dedicated BFD impacts traffic throughput as one CPU core is removed from data plane processing.

- To enable real-time BFD on SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and SRX650 devices, use the **set chassis realtime-ukern-thread** command.

Enabling real-time BFD does not impact data plane performance. Higher priority is given to the pfe process handling BFD in distributed mode. This is suitable for scenarios where the number of BFD sessions are less.

Table 4 on page 101 lists the BFD modes supported on SRX Series devices.

**Table 4: BFD Modes Supported on SRX Series Devices**

SRX Series Device	Default BFD Mode	Distributed BFD	Real-Time BFD	Dedicated Core
SRX100	Centralized	Configuration	Configuration (Optional)	Not supported
SRX110	Centralized	Configuration	Configuration (Optional)	Not supported

Table 4: BFD Modes Supported on SRX Series Devices (continued)

SRX Series Device	Default BFD Mode	Distributed BFD	Real-Time BFD	Dedicated Core
SRX210	Centralized	Configuration	Configuration (Optional)	Not supported
SRX220	Centralized	Configuration	Configuration (Optional)	Not supported
SRX240	Centralized	Configuration	Configuration	Configuration (Optional)
SRX300	Centralized	Configuration	Configuration (Optional)	Not supported
SRX320	Centralized	Configuration	Configuration (Optional)	Not supported
SRX340	Centralized	Configuration	Configuration	Configuration (Optional)
SRX345	Centralized	Configuration	Configuration	Configuration (Optional)
SRX550	Centralized	Configuration	Configuration	Configuration (Optional)
SRX550M	Centralized	Configuration	Configuration	Configuration (Optional)
SRX650	Centralized	Configuration	Configuration	Configuration (Optional)
SRX1500	Centralized	Configuration	Not supported	Configuration (Optional)
SRX4100	Centralized	Not supported	Not supported	Not supported
SRX4200	Centralized	Not supported	Not supported	Not supported
SRX5400	Centralized	Not supported	Not supported	Not supported
SRX5600	Centralized	Not supported	Not supported	Not supported
SRX5800	Centralized	Not supported	Not supported	Not supported

To determine if a BFD peer is running distributed BFD, run the **show bfd sessions extensive** command and look for **Remote is control-plane independent** in the command output.

For distributed BFD to work, you need to configure the lo0 interface with unit 0 and the appropriate family.

```
# set interfaces lo0 unit 0 family inet
# set interfaces lo0 unit 0 family inet6
# set interfaces lo0 unit 0 family mpls
```

This is true for the following types of BFD sessions:

- BFD over ae logical interfaces, both IPv4 and IPv6
- Multihop BFD, both IPv4 and IPv6

- BFD over VLAN interfaces in EX Series switches, both IPv4 and IPv6
- Virtual Circuit Connectivity Verification (VCCV) BFD (Layer 2 circuit, Layer 3 VPN, and VPLS) (MPLS)



**NOTE:** Starting in Junos OS Release 13.3R5, if you apply a firewall filter on a loopback interface for a multihop BFD session with a delegated anchor FPC, Junos OS does not execute this filter, because there is an implicit filter on all ingress FPCs to forward packets to the anchor FPC. Therefore, the firewall filter on the loopback interface is not applied on these packets. If you do not want these packets to be forwarded to the anchor FPC, you can configure the `no-delegate-processing` option.

For information about troubleshooting BFD, see [Juniper Networks Knowledge Base article 26746](#).



**NOTE:** Starting in Junos OS Release 13.3, the distribution of adjacency entry (the IP addresses of adjacent routers) and transmit entry (the IP address of transmitting routers) for a BFD session is asymmetric. This is because an adjacency entry that requires rules might or might not be distributed based on the redirect rule, and the distribution of transmit entries is *not* dependent on the redirect rule.

The term *redirect rule* here denotes the capability of an interface to send protocol redirect messages. See *Disabling the Transmission of Redirect Messages on an Interface*.

Release History Table

Release	Description
15.1X49-D100	Starting with Junos OS Release 15.1X49-D100, dedicated BFD is supported on SRX340, SRX345, and SRX1500 devices.
15.1X49-D100	Starting with Junos OS Release 15.1X49-D100, real-time BFD is supported on SRX300 and SRX320 devices.
13.3R5	Starting in Junos OS Release 13.3R5, if you apply a firewall filter on a loopback interface for a multihop BFD session with a delegated anchor FPC, Junos OS does not execute this filter, because there is an implicit filter on all ingress FPCs to forward packets to the anchor FPC.
13.3	Starting in Junos OS Release 13.3, the distribution of adjacency entry (the IP addresses of adjacent routers) and transmit entry (the IP address of transmitting routers) for a BFD session is asymmetric.

#### Related Documentation

- `show bfd session`
- [Understanding BFD for RIP on page 75](#)

- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 37](#)
- [Understanding BFD for BGP on page 48](#)
- [Understanding BFD for IS-IS on page 66](#)
- [Understanding BFD for OSPF on page 59](#)
- *Understanding EBGP Multihop*



## PART 1

# Configuration Statements and Operational Commands

- [Configuration Statements: Bidirectional Forwarding Detection on page 107](#)
- [Configuration Statements: Graceful Routing Engine Switchover on page 115](#)
- [Configuration Statements: Graceful Restart on page 117](#)
- [Configuration Statements: Nonstop Active Routing on page 137](#)
- [Configuration Statements: Nonstop Bridging on page 145](#)
- [Configuration Statements: Routing Engine and Switching Control Board Redundancy on page 147](#)
- [Configuration Statements: Unified ISSU on page 161](#)
- [Configuration Statements: VRRP on page 165](#)
- [Operational Commands on page 199](#)



## CHAPTER 4

# Configuration Statements: Bidirectional Forwarding Detection

- [authentication \(LAG\) on page 108](#)
- [bfd-liveness-detection \(LAG\) on page 109](#)
- [detection-time \(LAG\) on page 111](#)
- [traceoptions \(Protocols BFD\) on page 112](#)
- [transmit-interval \(LAG\) on page 114](#)

## authentication (LAG)

---

<b>Syntax</b>	<pre>authentication {     algorithm <i>algorithm-name</i>;     key-chain <i>key-chain-name</i>;     loose-check; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>aex</i> aggregated-ether-options bfd-liveness-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	Configure the authentication criteria of the BFD session for aggregated Ethernet interfaces.
<b>Options</b>	<p><b>algorithm <i>algorithm-name</i></b>—Specify the algorithm to be used to authenticate the BFD session. You can use one of the following algorithms for authentication:</p> <ul style="list-style-type: none"><li>• keyed-md5</li><li>• keyed-sha-1</li><li>• meticulous-keyed-md5</li><li>• meticulous-keyed-sha-1</li><li>• simple-password</li></ul> <p><b>key-chain <i>key-chain-name</i></b>—Specify the name that is associated with the security key for the BFD session. The name you specify must match one of the keychains configured in the <b>authentication-key-chains key-chain</b> statement at the [edit security] hierarchy level.</p> <p><b>loose-check</b>—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">bfd-liveness-detection on page 109</a></li><li>• <a href="#">detection-time on page 111</a></li><li>• <a href="#">transmit-interval on page 114</a></li><li>• <a href="#">Configuring Independent Micro BFD Sessions for LAG on page 85</a></li><li>• <a href="#">Example: Configuring Independent Micro BFD Sessions for LAG on page 91</a></li><li>• <a href="#">Understanding Independent Micro BFD Sessions for LAG on page 83</a></li></ul>

## bfd-liveness-detection (LAG)

**Syntax**

```

bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  holddown-interval milliseconds;
  local-address bfd-local-address;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  neighbor bfd-neighbor-address;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}

```

**Hierarchy Level** [edit interfaces *aex* aggregated-ether-options]

**Release Information** Statement introduced in Junos OS Release 13.3.

**Description** Configure Bidirectional Forwarding Detection (BFD) timers and authentication for aggregated Ethernet interfaces.

**Options** **holddown-interval *milliseconds***— Specify a time limit, in milliseconds, indicating the time that a BFD session remains up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

**Range:** 0 through 255,000

**Default:** 0

**local-address *bfd-local-address***— Specify the loopback address or the AE interface address of the source of the BFD session.



**NOTE:** Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD **local-address** against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

**minimum-interval *milliseconds***— Specify a minimum time interval after which the local routing device transmits a BFD packet and then expects to receive a reply from the BFD neighbor. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** statement.

**Range:** 1 through 255,000

**minimum-receive-interval *milliseconds***— Specify the minimum time interval after which the routing device expects to receive a reply from the BFD neighbor.

**Range:** 1 through 255,000

**multiplier *number***— Specify the number of BFD packets that were not received by the BFD neighbor before the originating interface is declared down.

**Range:** 1 through 255

**neighbor *bfd-neighbor-address***— Specify the loopback address or the AE interface address of a remote destination to send BFD packets.

**no-adaptation**— Disable the BFD adaptation. Include this statement if you do not want the BFD sessions to adapt to changing network conditions. We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**version**— Configure the BFD version to detect (BFD version 1) or autodetect (the BFD version).



**NOTE:** The version option is not supported on the QFX Series.

---

**Default:** automatic

The remaining statements are explained separately. See [CLI Explorer](#).

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">authentication on page 108</a></li><li>• <a href="#">detection-time on page 111</a></li><li>• <a href="#">transmit-interval on page 114</a></li><li>• <a href="#">Configuring Independent Micro BFD Sessions for LAG on page 85</a></li><li>• <a href="#">Example: Configuring Independent Micro BFD Sessions for LAG on page 91</a></li><li>• <a href="#">Understanding Independent Micro BFD Sessions for LAG on page 83</a></li></ul>
------------------------------	--

## detection-time (LAG)

---

<b>Syntax</b>	<pre>detection-time {     threshold <i>milliseconds</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces aex aggregated-ether-options bfd-liveness-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	Configure BFD timers for aggregated Ethernet interfaces.
<b>Options</b>	<b>threshold <i>milliseconds</i></b> — Specify the maximum time interval for detecting a BFD neighbor. If the transmit interval is greater than this value, the device triggers a trap.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">authentication on page 108</a></li><li>• <a href="#">bfd-liveness-detection on page 109</a></li><li>• <a href="#">transmit-interval on page 114</a></li><li>• <a href="#">Configuring Independent Micro BFD Sessions for LAG on page 85</a></li><li>• <a href="#">Example: Configuring Independent Micro BFD Sessions for LAG on page 91</a></li><li>• <a href="#">Understanding Independent Micro BFD Sessions for LAG on page 83</a></li></ul>

## traceoptions (Protocols BFD)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>name</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>
<b>Hierarchy Level</b>	[edit protocols bfd]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. issu flag for BFD added in Junos OS Release 9.1.
<b>Description</b>	<p>Define tracing operations that track unified in-service software upgrade (ISSU) functionality in the router.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>
<b>Default</b>	If you do not include this statement, no global tracing operations are performed.
<b>Options</b>	<p><b>disable</b>—(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 <b>all</b>.</p> <p><b>file <i>name</i></b>—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 <b>/var/log</b>. We recommend that you place global routing protocol tracing output in the file <b>routing-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. The tracing options are as follows:</p> <ul style="list-style-type: none"><li>• <b>adjacency</b>—Trace adjacency messages.</li><li>• <b>all</b>—Trace everything.</li><li>• <b>error</b>—Trace all errors.</li><li>• <b>events</b>—Trace all events.</li><li>• <b>issu</b>—Trace ISSU packet activity.</li></ul>



- **nsr-packet**—Trace packet activity of NSR.
- **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.
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

**no-world-readable**—Restrict users 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.

**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.

<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Managing and Tracing BFD Sessions During Unified ISSU Procedures</i></li> </ul>

## transmit-interval (LAG)

---

<b>Syntax</b>	<pre>transmit-interval {   minimum-interval <i>milliseconds</i>;   threshold <i>milliseconds</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces aex aggregated-ether-options bfd-liveness-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	Configure the minimum interval and the threshold for transmission of BFD packets for aggregated Ethernet interfaces.
<b>Options</b>	<p><b>minimum-interval <i>milliseconds</i></b>— Specify the minimum time interval between two transmissions of packets. <b>Range:</b> 1 through 255,000</p> <p><b>threshold <i>milliseconds</i></b>— Specify the maximum interval between transmission of packets. If the transmit interval is greater than this value, the device triggers a trap.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">authentication on page 108</a></li><li>• <a href="#">bfd-liveness-detection on page 109</a></li><li>• <a href="#">detection-time on page 111</a></li><li>• <a href="#">Configuring Independent Micro BFD Sessions for LAG on page 85</a></li><li>• <a href="#">Example: Configuring Independent Micro BFD Sessions for LAG on page 91</a></li><li>• <a href="#">Understanding Independent Micro BFD Sessions for LAG on page 83</a></li></ul>

## CHAPTER 5

# Configuration Statements: Graceful Routing Engine Switchover

- [graceful-switchover](#) on page 115

### [graceful-switchover](#)

---

<b>Syntax</b>	<code>graceful-switchover;</code>
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Graceful Routing Engine Switchover</i></li></ul>



## CHAPTER 6


# Configuration Statements: Graceful Restart

- [disable](#) on page 118
- [graceful-restart \(Enabling Globally\)](#) on page 119
- [graceful-restart \(Multicast Snooping\)](#) on page 120
- [helper-disable \(Multiple Protocols\)](#) on page 121
- [helper-disable \(OSPF\)](#) on page 122
- [maximum-helper-recovery-time](#) on page 123
- [maximum-helper-restart-time \(RSVP\)](#) on page 124
- [maximum-neighbor-reconnect-time](#) on page 125
- [maximum-neighbor-recovery-time](#) on page 126
- [no-strict-lsa-checking](#) on page 127
- [notify-duration](#) on page 128
- [not-on-disk-underperform](#) on page 129
- [reconnect-time](#) on page 130
- [recovery-time](#) on page 131
- [restart-duration](#) on page 132
- [restart-time \(BGP Graceful Restart\)](#) on page 133
- [stale-routes-time](#) on page 134
- [traceoptions \(Protocols\)](#) on page 135

## disable

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (bgp   isis   ldp   ospf   ospf3   pim   rip   ripng   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (bgp   ldp   ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit protocols (bgp   isis   isis   ospf   ospf3   ldp   pim   rip   ripng   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit protocols bgp group <i>group-name</i> <a href="#">graceful-restart</a>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>ip-address</i> <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (bgp   ldp   ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">graceful-restart</a>],</p> <p>[edit routing-options <a href="#">graceful-restart</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Disable graceful restart.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Enabling Graceful Restart</i></li> <li>• <i>Configuring Routing Protocols Graceful Restart</i></li> <li>• <i>Configuring Graceful Restart for MPLS-Related Protocols</i></li> <li>• <i>Configuring VPN Graceful Restart</i></li> <li>• <i>Configuring Logical System Graceful Restart</i></li> <li>• <i>Graceful Restart Configuration Statements</i></li> <li>• <i>Configuring Graceful Restart for QFabric Systems</i></li> </ul>

## graceful-restart (Enabling Globally)

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	<pre> [edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options] </pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>You configure the graceful restart routing option globally to enable the feature, but not to enable graceful restart for all routing protocols in a routing instance. Because all routing protocols are not usually run on every routing instance, you must also configure graceful restart for individual routing protocols running on a routing instance, including the main routing instance. You can, optionally, modify the global settings at the individual protocol level.</p>
	<div>  <p><b>NOTE:</b></p> <ul style="list-style-type: none"> <li>For VPNs, the <code>graceful-restart</code> statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.</li> <li>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.</li> <li>LDP sessions flap when <code>graceful-restart</code> configurations change.</li> </ul> </div>
<b>Default</b>	Graceful restart is disabled by default.
<b>Options</b>	The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .

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

**Related Documentation**

- *Enabling Graceful Restart*
- *Configuring Routing Protocols Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Configuring Graceful Restart for QFabric Systems*

---

## graceful-restart (Multicast Snooping)

---

**Syntax** graceful-restart {  
    disable;  
    restart-duration *seconds*;  
}

**Hierarchy Level** [edit multicast-snooping-options]

**Release Information** Statement introduced in Junos OS Release 9.2.

**Description** Establish the graceful restart duration for multicast snooping. You can set this value between 0 and 300 seconds. If you set the duration to 0, graceful restart is effectively disabled. Set this value slightly larger than the IGMP query response interval.

**Default** 180 seconds

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

**Related Documentation**


- *Example: Configuring Multicast Snooping*
- *query-response-interval (Bridge Domains)*



## helper-disable (Multiple Protocols)

<b>Syntax</b>	helper-disable;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (isis   ldp   ospf   ospf3   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ldp   ospf   ospf3) <a href="#">graceful-restart</a>],</p> <p>[edit protocols (isis   ldp   ospf   ospf3   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ldp   ospf   ospf3) <a href="#">graceful-restart</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	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.
<b>Default</b>	Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Routing Protocols Graceful Restart</i></li> <li>• <i>Configuring Graceful Restart for MPLS-Related Protocols</i></li> </ul>

## helper-disable (OSPF)

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

## maximum-helper-recovery-time

---

<b>Syntax</b>	<code>maximum-helper-recovery-time seconds;</code>
<b>Hierarchy Level</b>	[edit protocols rsvp <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	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.
<b>Options</b>	<b>seconds</b> —Length of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart. <b>Range:</b> 1 through 3600 <b>Default:</b> 180
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Graceful Restart Options for RSVP, CCC, and TCC</i></li> <li>• <a href="#">maximum-helper-restart-time (RSVP) on page 124</a></li> </ul>

## maximum-helper-restart-time (RSVP)

---

<b>Syntax</b>	<code>maximum-helper-restart-time seconds;</code>
<b>Hierarchy Level</b>	[edit protocols rsvp <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	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.
<b>Options</b>	<b>seconds</b> —The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. <b>Range:</b> 1 through 1800 <b>Default:</b> 60
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Graceful Restart Options for RSVP, CCC, and TCC</i></li><li>• <a href="#">maximum-helper-recovery-time on page 123</a></li></ul>

---

## maximum-neighbor-reconnect-time

---

<b>Syntax</b>	maximum-neighbor-reconnect-time <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit protocols ldp <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> protocols ldp <a href="#">graceful-restart</a> ], [edit routing-instances <i>routing-instance-name</i> protocols ldp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1.
<b>Description</b>	Specify the maximum length of time allowed to reestablish connection from a restarting neighbor.
<b>Options</b>	<b><i>seconds</i></b> —Maximum time allowed for reconnection. <b>Range:</b> 30 through 300
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Graceful Restart Options for LDP</a></li></ul>

## maximum-neighbor-recovery-time

---

<b>Syntax</b>	<code>maximum-neighbor-recovery-time seconds;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement changed from <b>maximum-recovery-time</b> to <b>maximum-neighbor-recovery-time</b> in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
<b>Options</b>	<b>seconds</b> —Configure the maximum recovery time, in seconds. <b>Range:</b> 120 through 1800 seconds <b>Default:</b> 140 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Recovery Time and Maximum Recovery Time</i></li><li>• <i>Configuring Graceful Restart Options for LDP</i></li><li>• <a href="#">no-strict-lsa-checking on page 127</a></li><li>• <a href="#">recovery-time on page 131</a></li></ul>

## no-strict-lsa-checking

---

<b>Syntax</b>	no-strict-lsa-checking;
<b>Hierarchy Level</b>	[edit protocols (ospf   ospf3) <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router or switch.
<b>Default</b>	By default, LSA checking is enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Graceful Restart Options for OSPF and OSPFv3</i></li> <li>• <i>Configuring Graceful Restart for QFabric Systems</i></li> <li>• <a href="#">maximum-neighbor-recovery-time on page 126</a></li> <li>• <a href="#">recovery-time on page 131</a></li> </ul>

## notify-duration

---

<b>Syntax</b>	<code>notify-duration <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit protocols (ospf   ospf3) <a href="#">graceful-restart</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a>],</code> <code>[edit routing-instances <i>instance-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify the length of time the router or switch notifies helper OSPF routers that it has completed graceful restart.
<b>Options</b>	<b><i>seconds</i></b> —Length of time in the router notifies helper OSPF routers that it has completed graceful restart. <b>Range:</b> 1 through 3600 <b>Default:</b> 30
<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Graceful Restart Options for OSPF and OSPFv3</i></li><li>• <i>Configuring Graceful Restart for QFabric Systems</i></li><li>• <a href="#">restart-duration on page 132</a></li></ul>



## not-on-disk-underperform

---

**Syntax** not-on-disk-underperform;

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

**Release Information** Statement introduced in Junos OS Release 13.3R6.

**Description** Prevent gstatd from causing failovers in dual Routing Engines set for graceful Routing Engine switchover (GRES). The gstatd log message is still generated. This is an optional configuration.



**NOTE:** Configure the `disk-write-threshold` and `disk-read-threshold` statements to customize the gstatd timeout threshold.

---

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

**Related Documentation**

- *Preventing Graceful Routing Engine Switchover in the Case of Slow Disks*

## reconnect-time

---

<b>Syntax</b>	<code>reconnect-time seconds;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
<b>Options</b>	<b>seconds</b> —Time required for reconnection. <b>Range:</b> 30 through 300 <b>Default:</b> 60 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring LDP Graceful Restart on MPLS Applications Feature Guide</i></li><li>• <i>Configuring Graceful Restart Options for LDP</i></li></ul>

## recovery-time

---

<b>Syntax</b>	<code>recovery-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp <a href="#">graceful-restart</a> ], [edit protocols ldp <a href="#">graceful-restart</a> ], [edit routing-instances <i>routing-instance-name</i> protocols ldp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the length of time a router or switch waits for Label Distribution Protocol (LDP) neighbors to assist it with a graceful restart.
<b>Options</b>	<b>seconds</b> —Time the router waits for LDP to restart gracefully. <b>Range:</b> 120 through 1800 <b>Default:</b> 160
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for LDP</a></li> <li>• <a href="#">maximum-neighbor-recovery-time on page 126</a></li> <li>• <a href="#">no-strict-lsa-checking on page 127</a></li> </ul>

## restart-duration

<b>Syntax</b>	<code>restart-duration seconds;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (isis   ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit protocols (esis   isis   ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit routing-options <a href="#">graceful-restart</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<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>
<b>Options</b>	<p><b><i>seconds</i></b>—Time for the graceful restart period.</p> <p><b>Range:</b></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"> <li>• <b>[edit routing-options graceful-restart]</b> (global setting)—120 through 900</li> <li>• ES-IS—30 through 300</li> <li>• IS-IS—30 through 300</li> <li>• OSPF/OSPFv3—1 through 3600</li> <li>• PIM—30 through 300</li> </ul> <p><b>Default:</b></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"> <li>• <b>[edit routing-options graceful-restart]</b> (global setting)—300</li> <li>• ES-IS—180</li> <li>• IS-IS—210</li> <li>• OSPF/OSPFv3—180</li> <li>• PIM—60</li> </ul>

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

**Related Documentation**

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

## restart-time (BGP Graceful Restart)

**Syntax** restart-time *seconds*;

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

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

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

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

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

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

**Related Documentation**

- *Configuring Graceful Restart Options for BGP*
- *Configuring Graceful Restart Options for RIP and RIPng*
- *Configuring Graceful Restart for QFabric Systems*
- [stale-routes-time on page 134](#)

## stale-routes-time

---

<b>Syntax</b>	<code>stale-routes-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-routing-name</i> protocols bgp <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a> ], [edit protocols bgp <a href="#">graceful-restart</a> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
<b>Description</b>	Specify the maximum time that stale routes are kept during a restart. The <b>stale-routes-time</b> statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.
<b>Options</b>	<b>seconds</b> —Time the router device waits to receive messages from restarting neighbors before declaring them down. <b>Range:</b> 1 through 600 seconds <b>Default:</b> 300 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Graceful Restart Options for BGP</i></li><li>• <i>Configuring Graceful Restart for QFabric Systems</i></li><li>• <a href="#">restart-time (BGP Graceful Restart) on page 133</a></li></ul>

## traceoptions (Protocols)

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

***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 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.

<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Tracking Graceful Restart Events</i></li></ul>



## CHAPTER 7

# Configuration Statements: Nonstop Active Routing

- [nonstop-routing on page 138](#)
- [switchover-on-routing-crash on page 139](#)
- [synchronize on page 140](#)
- [traceoptions on page 142](#)

## nonstop-routing

---

**Syntax**    nonstop-routing;

**Hierarchy Level**    [edit routing-options]



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

**Release Information**    Statement introduced in Junos OS Release 8.4.  
Statement introduced in Junos OS Release 10.4 for EX Series switches.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.  
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches

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

**Default**    disabled

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

**Related Documentation**    • *Configuring Nonstop Active Routing*

---

## switchover-on-routing-crash

---

<b>Syntax</b>	switchover-on-routing-crash;
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3 for M Series, MX Series, T Series, TX Matrix, PTX Series, EX Series, QFX Series.
<b>Description</b>	Prevent loss of traffic in the case of NSR being configured. With the <b>switchover-on-routing-crash</b> configuration statement enabled, when rpd on the master Routing Engine crashes with NSR configured, the Routing Engine will switch over immediately to the backup Routing Engine to preserve protocol state and adjacencies. Prior to having this statement, if NSR was configured and rpd on the master Routing Engine crashed, it would cause network impact (protocol neighbor and adjacency drops and traffic loss).
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Nonstop Active Routing</i></li></ul>

## synchronize

---

<b>Syntax</b>	synchronize;
<b>Hierarchy Level</b>	[edit system commit]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 10.4 for EX Series switches.
<b>Description</b>	For devices with multiple Routing Engines only. Configure the <b>commit</b> command to automatically perform a <b>commit synchronize</b> action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the <b>commit</b> command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.



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



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

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

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

Engine on the TX Matrix Plus router distributes the configuration to the backup Routing Engine on each LCC.

In EX Series Virtual Chassis configurations:

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

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

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

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

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

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

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

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

**Related**    • *Synchronizing the Routing Engine Configuration*

**Documentation**    • *Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically*

## traceoptions

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; } </pre>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>nsr-synchronization</b> flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p><b>nsr-synchronization</b> and <b>nsr-packet</b> 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><b>nsr-synchronization</b> flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p><b>nsr-synchronization</b> flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p><b>nsr-synchronization</b> flag for PIM added in Junos OS Release 9.3.</p> <p><b>nsr-synchronization</b> 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><b>nsr-synchronization</b> flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<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 <b>flag</b> statements.</p>
<b>Default</b>	If you do not include this statement, no global tracing operations are performed.
<b>Options</b>	<p><b>Values:</b></p> <p><b>disable</b>—(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 <b>all</b>.</p> <p><b>file <i>filename</i></b>—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 <b>/var/log</b>. We recommend that you place global routing protocol tracing output in the file <b>routing-log</b>.</p>

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

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

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

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

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

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

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

**Default:** 128 KB

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

<b>Required Privilege</b>	routing and trace—To view this statement in the configuration.
<b>Level</b>	routing-control and trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Tracing Global Routing Protocol Operations</i></li></ul>



## CHAPTER 8

# Configuration Statements: Nonstop Bridging

- [nonstop-bridging on page 145](#)

### nonstop-bridging

---

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



## CHAPTER 9

# Configuration Statements: Routing Engine and Switching Control Board Redundancy

- [cfeb](#) on page 148
- [description \(Chassis Redundancy\)](#) on page 148
- [failover \(Chassis\)](#) on page 149
- [failover \(System Process\)](#) on page 150
- [feb \(Creating a Redundancy Group\)](#) on page 151
- [feb \(Assigning a FEB to a Redundancy Group\)](#) on page 152
- [keepalive-time](#) on page 153
- [no-auto-failover](#) on page 154
- [on-disk-failure \(Chassis Redundancy Failover\)](#) on page 154
- [on-loss-of-keepalives](#) on page 155
- [redundancy](#) on page 156
- [redundancy-group](#) on page 157
- [routing-engine \(Chassis Redundancy\)](#) on page 158
- [sfm \(Chassis Redundancy\)](#) on page 159
- [ssb](#) on page 160

## cfeb

---

<b>Syntax</b>	<code>cfeb slot-number</code> (always   preferred);
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	On M10i routers only, configure which Compact Forwarding Engine Board (CFEB) is the master and which is the backup.
<b>Default</b>	By default, the CFEB in slot 0 is the master and the CFEB in slot 1 is the backup.
<b>Options</b>	<b>slot-number</b> —Specify which slot is the master and which is the backup.  <b>always</b> —Define this CFEB as the sole device.  <b>preferred</b> —Define this CFEB as the preferred device of at least two.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring CFEB Redundancy on the M10i Router on page 29</a></li></ul>

## description (Chassis Redundancy)

---

<b>Syntax</b>	<code>description description</code> ;
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy</a> <a href="#">feb</a> <a href="#">redundancy-group</a> <i>group-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Provide a description of the FEB redundancy group.
<b>Options</b>	<b>description</b> —Provide a description for the FEB redundancy group.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring FEB Redundancy on the M120 Router on page 30</a></li></ul>

## failover (Chassis)

---

<b>Syntax</b>	<pre>failover {   on-disk-failure;   on-loss-of-keepalives;   on-re-to-fpc-stale; }</pre>
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>on-re-to-fpc-stale</b> option introduced in Junos OS Release 15.2 on the MX240, MX480, MX960, MX2010, and MX2020.
<b>Description</b>	<p>Specify conditions on the master Routing Engine that cause the backup router to take mastership.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>On Detection of a Hard Disk Error on the Master Routing Engine</i></li></ul>

## failover (System Process)

---

<b>Syntax</b>	<code>failover (alternate-media   other-routing-engine);</code>
<b>Hierarchy Level</b>	[edit system processes <i>process-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the router to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
<b>Options</b>	<p><b><i>process-name</i></b>—Junos OS process name. Some of the processes that support the <b>failover</b> statement are <b>bootp</b>, <b>chassis-control</b>, <b>craft-control</b>, <b>ethernet-connectivity-fault-management</b>, <b>init</b>, <b>interface-control</b>, <b>neighbor-liveness</b>, <b>pfe</b>, <b>redundancy-interface-process</b>, <b>routing</b>, <b>smg-service</b>, and <b>vrrp</b>.</p> <p><b>alternate-media</b>—Use the Junos OS image on alternate media during the reboot.</p> <p><b>other-routing-engine</b>—On routers with dual Routing Engines, use the Junos OS image on the other Routing Engine during the reboot. That Routing Engine assumes mastership; in the usual configuration, the other Routing Engine is the designated backup Routing Engine.</p>
<b>Required Privilege Level</b>	<p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>When a Software Process Fails</i></li><li>• <i>processes</i></li></ul>

## feb (Creating a Redundancy Group)

---

**Syntax**    feb {  
              **redundancy-group** *group-name* {  
                  **description** *description*;  
                  **feb** *slot-number* (backup | primary);  
                  **no-auto-failover**;  
              }  
          }

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

**Release Information**    Statement introduced in Junos OS Release 8.2.

**Description**    On M120 routers only, configure a Forwarding Engine Board (FEB) redundancy group.

**Options**    The remaining statements are described separately.

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

**Related Documentation**    • [Configuring FEB Redundancy on the M120 Router on page 30](#)

## feb (Assigning a FEB to a Redundancy Group)

---

<b>Syntax</b>	<code>feb slot-number (backup   primary);</code>
<b>Hierarchy Level</b>	[edit chassis <b>redundancy feb redundancy-group</b> group-name]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	On M120 routers only, configure a Forwarding Engine Board (FEB) as part of a FEB redundancy group.
<b>Options</b>	<p><b>slot-number</b>—Slot number of the FEB. The range of values is from 0 to 5.</p> <p><b>backup</b>—(Optional) For each redundancy group, you must configure exactly one backup FEB.</p> <p><b>primary</b>—(Optional) For each redundancy group, you can optionally configure one primary FEB.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring FEB Redundancy on the M120 Router on page 30</a></li></ul>



## keepalive-time

---

<b>Syntax</b>	<code>keepalive-time <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit chassis <a href="#">redundancy</a>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the time period that must elapse before the backup router takes mastership when it detects loss of the keepalive signal.
<b>Default</b>	<p>The <b>on-loss-of-keepalives</b> statement at the <code>[edit chassis redundancy failover]</code> hierarchy level must be included for failover to occur.</p> <p>When the <b>on-loss-of-keepalives</b> statement is included and graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the <b>on-loss-of-keepalives</b> statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds (4 seconds on M20 routers). You cannot manually reset the keepalive time.</p>
<b>Options</b>	<b><i>seconds</i></b> —Time before the backup router takes mastership when it detects loss of the keepalive signal. The range of values is 2 through 10,000.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>On Detection of a Loss of Keepalive Signal from the Master Routing Engine</i></li><li>• <a href="#">failover (Chassis) on page 149</a></li><li>• <a href="#">on-loss-of-keepalives on page 155</a></li></ul>

## no-auto-failover

---

<b>Syntax</b>	no-auto-failover;
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy feb redundancy-group group-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Disable automatic failover to a backup FEB when an active FEB in a redundancy group fails.
<b>Default</b>	Automatic failover is enabled by default.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring FEB Redundancy on the M120 Router on page 30</a></li></ul>

## on-disk-failure (Chassis Redundancy Failover)

---

<b>Syntax</b>	on-disk-failure;
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy failover</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Instruct the backup router to take mastership if it detects hard disk errors on the master Routing Engine.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>On Detection of a Hard Disk Error on the Master Routing Engine</i></li></ul>

## on-loss-of-keepalives

---

<b>Syntax</b>	on-loss-of-keepalives;
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy failover</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Instruct the backup router to take mastership if it detects a loss of keepalive signal from the master Routing Engine.
<b>Default</b>	<p>The <b>on-loss-of-keepalives</b> statement must be included at the <b>[edit chassis redundancy failover]</b> hierarchy level for failover to occur.</p> <p>When the <b>on-loss-of-keepalives</b> statement is included but graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the <b>on-loss-of-keepalives</b> statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds (4 seconds on M20 routers) . The keepalive time is not configurable.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>On Detection of a Loss of Keepalive Signal from the Master Routing Engine</i></li><li>• <a href="#">keepalive-time on page 153</a></li></ul>

## redundancy

<b>Syntax</b>	<pre> redundancy {   cfcb slot (always   preferred);   failover {     on-disk-failure;     on-loss-of-keepalives;     on-re-to-fpc-stale;   }   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); } </pre>
<b>Hierarchy Level</b>	[edit chassis]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure redundancy options.
<b>Options</b>	The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Routing Engine Redundancy</a></li> <li>• <a href="#">Configuring CFEB Redundancy on the M10i Router on page 29</a></li> <li>• <a href="#">Configuring FEB Redundancy on the M120 Router on page 30</a></li> <li>• <a href="#">Configuring SFM Redundancy on M40e and M160 Routers on page 33</a></li> <li>• <a href="#">Configuring SSB Redundancy on the M20 Router on page 34</a></li> </ul>

## redundancy-group

---

<b>Syntax</b>	<pre>redundancy-group <i>group-name</i> {   <i>description</i> <i>description</i>;   feb <i>slot-number</i> (backup   primary);   no-auto-failover; }</pre>
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy feb</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	On M120 routers only, configure a Forwarding Engine Board (FEB) redundancy group.
<b>Options</b>	<p><i>group-name</i> is the unique name for the redundancy group. The maximum length is 39 alphanumeric characters.</p> <p>Other statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring FEB Redundancy on the M120 Router on page 30</a></li></ul>

## routing-engine (Chassis Redundancy)

---

<b>Syntax</b>	<code>routing-engine slot-number (backup   disabled   master);</code>
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure Routing Engine redundancy.
<b>Default</b>	By default, the Routing Engine in slot 0 is the master Routing Engine and the Routing Engine in slot 1 is the backup Routing Engine.
<b>Options</b>	<p><i>slot-number</i>—Specify the slot number (0 or 1).</p> <p>Set the function of the Routing Engine for the specified slot:</p> <ul style="list-style-type: none"><li>• <b>master</b>—Routing Engine in the specified slot is the master.</li><li>• <b>backup</b>—Routing Engine in the specified slot is the backup.</li><li>• <b>disabled</b>—Routing Engine in the specified slot is disabled.</li></ul>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Routing Engine Redundancy</i></li></ul>

## sfm (Chassis Redundancy)

---

<b>Syntax</b>	<code>sfm slot-number (always   preferred);</code>
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	On M40e and M160 routers, configure which Switching and Forwarding Module (SFM) is the master and which is the backup.
<b>Default</b>	By default, the SFM in slot 0 is the master and the SFM in slot 1 is the backup.
<b>Options</b>	<p><b>slot-number</b>—Specify which slot is the master and which is the backup. On the M40e router, <b>slot-number</b> can be 0 or 1. On the M160 router, <b>slot-number</b> can be 0 through 3.</p> <p><b>always</b>—Define this SFM as the sole device.</p> <p><b>preferred</b>—Define this SFM as the preferred device of at least two.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SFM Redundancy on M40e and M160 Routers on page 33</a></li></ul>

## ssb

---

<b>Syntax</b>	<code>ssb slot-number (always   preferred);</code>
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	On M20 routers, configure which System and Switch Board (SSB) is the master and which is the backup.
<b>Default</b>	By default, the SSB in slot 0 is the master and the SSB in slot 1 is the backup.
<b>Options</b>	<p><b>slot-number</b>—Specify which slot is the master and which is the backup.</p> <p><b>always</b>—Define this SSB as the sole device.</p> <p><b>preferred</b>—Define this SSB as the preferred device of at least two.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SSB Redundancy on the M20 Router on page 34</a></li></ul>



## CHAPTER 10

# Configuration Statements: Unified ISSU

- [no-issu-timer-negotiation](#) on page 161
- [traceoptions \(Protocols BFD\)](#) on page 162

### no-issu-timer-negotiation

<b>Syntax</b>	no-issu-timer-negotiation;
<b>Hierarchy Level</b>	[edit protocols bfd], [edit logical-systems <i>logical-system-name</i> protocols bfd], [edit routing-instances <i>routing-instance-name</i> protocols bfd]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 13.2 for PTX5000 routers.
<b>Description</b>	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.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Managing and Tracing BFD Sessions During Unified ISSU Procedures</i></li><li>• <i>Junos OS Routing Protocols Library</i></li></ul>

## traceoptions (Protocols BFD)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>name</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>
<b>Hierarchy Level</b>	[edit protocols bfd]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. issu flag for BFD added in Junos OS Release 9.1.
<b>Description</b>	<p>Define tracing operations that track unified in-service software upgrade (ISSU) functionality in the router.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>
<b>Default</b>	If you do not include this statement, no global tracing operations are performed.
<b>Options</b>	<p><b>disable</b>—(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 <b>all</b>.</p> <p><b>file <i>name</i></b>—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 <b>/var/log</b>. We recommend that you place global routing protocol tracing output in the file <b>routing-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. The tracing options are as follows:</p> <ul style="list-style-type: none"><li>• <b>adjacency</b>—Trace adjacency messages.</li><li>• <b>all</b>—Trace everything.</li><li>• <b>error</b>—Trace all errors.</li><li>• <b>events</b>—Trace all events.</li><li>• <b>issu</b>—Trace ISSU packet activity.</li></ul>

- **nsr-packet**—Trace packet activity of NSR.
- **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.
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

**no-world-readable**—Restrict users 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.

**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.

<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Managing and Tracing BFD Sessions During Unified ISSU Procedures</i></li> </ul>




## CHAPTER 11

# Configuration Statements: VRRP


- [accept-data on page 167](#)
- [advertise-interval on page 168](#)
- [asymmetric-hold-time on page 169](#)
- [authentication-key on page 170](#)
- [authentication-type on page 171](#)
- [bandwidth-threshold on page 172](#)
- [delegate-processing \(VRRP\) on page 173](#)
- [fast-interval on page 174](#)
- [global-advertisements-threshold on page 175](#)
- [hold-time \(VRRP\) on page 176](#)
- [inherit-advertisement-interval on page 177](#)
- [inet6-advertise-interval on page 178](#)
- [interface on page 179](#)
- [preempt \(VRRP\) on page 180](#)
- [priority \(Protocols VRRP\) on page 181](#)
- [priority-cost \(VRRP\) on page 182](#)
- [priority-hold-time on page 183](#)
- [route \(Interfaces\) on page 184](#)
- [skew-timer-disable on page 185](#)
- [startup-silent-period on page 186](#)
- [traceoptions \(Protocols VRRP\) on page 187](#)
- [track \(VRRP\) on page 189](#)
- [version-3 on page 190](#)
- [virtual-address on page 191](#)
- [virtual-inet6-address on page 192](#)
- [virtual-link-local-address on page 193](#)
- [vrrp-group on page 194](#)

- [vrrp-inet6-group](#) on page 196
- [vrrp-inherit-from](#) on page 197

## accept-data

<b>Syntax</b>	(accept-data   no-accept-data);
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
<b>Description</b>	<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"> <li>• <b>accept-data</b>—Enable the master router to accept all packets destined for the virtual IP address.</li> <li>• <b>no-accept-data</b>—Prevent the master router from accepting packets other than the ARP packets destined for the virtual IP address.</li> </ul>
<b>Default</b>	<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><b>NOTE:</b></p> <ul style="list-style-type: none"> <li>• If you want to restrict the incoming IP packets to ICMP packets only, you must configure firewall filters to accept only ICMP packets.</li> <li>• If you include the <b>accept-data</b> 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>).</li> </ul> </div>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group</i></li> </ul>

## advertise-interval

<b>Syntax</b>	<code>advertise-interval seconds;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <i>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> <b>vrrp-group</b> <i>group-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
<div>  <p><b>NOTE:</b> 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>	
<b>Options</b>	<p><b>seconds</b>—Interval between advertisement packets.</p> <p><b>Range:</b> 1 through 255 seconds</p> <p><b>Default:</b> 1 second</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Advertisement Interval for the VRRP Master Router</a></li> <li>• <a href="#">fast-interval on page 174</a></li> <li>• <a href="#">inet6-advertise-interval on page 178</a></li> <li>• <a href="#">version-3 on page 190</a></li> </ul>




## asymmetric-hold-time

---

<b>Syntax</b>	asymmetric-hold-time;
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5.
<b>Description</b>	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.
<b>Default</b>	asymmetric-hold-time is disabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Asymmetric Hold Time for VRRP Routers</i></li></ul>

## authentication-key

<b>Syntax</b>	<code>authentication-key key;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <i>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> <b>vrrp-group</b> <i>group-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 18.1R1 for the SRX Series devices.</p>
<b>Description</b>	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the <b>authentication-type</b> statement.</p> <p>All devices in the VRRP group must use the same authentication scheme and password.</p>
<div>  <p><b>NOTE:</b> When VRRPv3 is enabled, the <b>authentication-type</b> and <b>authentication-key</b> statements cannot be configured for any VRRP groups.</p> </div>	
<b>Options</b>	<p><b>key</b>—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 (" ").</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring VRRP Authentication (IPv4 Only)</a></li> <li>• <a href="#">authentication-type on page 171</a></li> <li>• <a href="#">version-3 on page 190</a></li> <li>• <a href="#">Understanding VRRP on SRX Series Devices</a></li> <li>• <a href="#">Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces</a></li> </ul>

## authentication-type

<b>Syntax</b>	<code>authentication-type (md5   simple);</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <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> <b>vrrp-group</b> <i>group-id</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series. Statement introduced in Junos OS Release 18.1R1 for the SRX Series devices.
<b>Description</b>	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 <b>authentication-key</b> statement.  All devices in the VRRP group must use the same authentication scheme and password.



**NOTE:** When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.

<b>Options</b>	<b>authentication</b> —Authentication scheme: <ul style="list-style-type: none"> <li><b>simple</b>—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure.</li> <li><b>md5</b>—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.</li> </ul> <p><b>Default:</b> none (no authentication is performed).</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring VRRP Authentication (IPv4 Only)</i></li> <li><a href="#">authentication-key on page 170</a></li> <li><a href="#">version-3 on page 190</a></li> <li><i>Understanding VRRP on SRX Series Devices</i></li> <li><i>Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces</i></li> </ul>


## bandwidth-threshold

<b>Syntax</b>	<code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>
<b>Hierarchy Level</b>	<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 interface 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 track interface 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 track interface 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 track interface interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
<b>Description</b>	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
<b>Options</b>	<p><b><i>bits-per-second</i></b>—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><b>Range:</b> 1 through 100000000000000 bits per second</p> <p><b><i>priority-cost priority</i></b>—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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring a Logical Interface to Be Tracked for a VRRP Group</li> </ul>


## delegate-processing (VRRP)

<b>Syntax</b>	<pre>delegate-processing {   ae-irb; }</pre>
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.6.</p> <p><b>ae-irb</b> option introduced in Junos OS Release 15.1.</p>
<b>Description</b>	<p>Configure the distributed periodic packet management process (ppmd) to send Virtual Router Redundancy Protocol (VRRP) advertisements .</p> <p>Using a hash logic based on iflIndex, the vrrp group ID, and the IP version, select one of the Flexible OIC Concentrators (FPCs) for distribution. The selected FPC is called the <i>anchor FPC</i>. All transmit instances and receive instances are from and to the anchor FPC. The anchor FPC is static, and VRRP is not guaranteed to get distributed to all available FPCs uniformly for all VRRP sessions.</p>
<b>Options</b>	<p><b>ae-irb</b>—Enable distributed ppmmd for VRRP over aggregated Ethernet and integrated routing and bridging (IRB) interfaces.</p> <p>Using the <b>ae-irb</b> option is only for MPC line cards. Using the <b>ae-irb</b> option requires use of the enhanced-ip mode.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Enabling the Distributed Periodic Packet Management Process for VRRP</i></li> </ul>

## fast-interval

<b>Syntax</b>	<code>fast-interval milliseconds;</code>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
<b>Description</b>	<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>
<b>Options</b>	<p><i>milliseconds</i>—Interval between advertisement packets.</p> <p><b>Range:</b> 10 through 40,950 milliseconds (range extended from 100–999 to 10–40,950 in Junos OS Release 12.2).</p>
<div>  <p><b>NOTE:</b> 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>	
<b>Default:</b> 1 second	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Advertisement Interval for the VRRP Master Router</a></li> <li>• <a href="#">advertise-interval on page 168</a></li> <li>• <a href="#">advertise-interval on page 168</a></li> <li>• <a href="#">inet6-advertise-interval on page 178</a></li> <li>• <a href="#">version-3 on page 190</a></li> </ul>

## global-advertisements-threshold

<b>Syntax</b>	<code>global-advertisements-threshold <i>advertisement-value</i>;</code>
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	Configure the number of fast advertisements that can be missed by a backup router before the master router is declared as down.
<div>  <b>NOTE:</b> <ul style="list-style-type: none"> <li>The advertisement value configured using the <code>global-advertisements-threshold</code> statement is applicable to all the Virtual Router Redundancy Protocol (VRRP) groups in the system.</li> <li>Setting the advertisement value of the <code>global-advertisements-threshold</code> 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 <code>global-advertisements-threshold</code> value to 1.</li> <li>Changing the advertisement value of the <code>global-advertisements-threshold</code> 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 <code>global-advertisements-threshold</code> statement during runtime.</li> </ul> </div>	
<b>Options</b>	<p><b><i>advertisement-value</i></b>—Number of VRRP advertisements missed before the master router is declared as down.</p> <p><b>Range:</b> 1 through 15</p> <p><b>Default:</b> 3</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Improving the Convergence Time for VRRP</i></li> <li><i>Configuring VRRP to Improve Convergence Time</i></li> </ul>

## hold-time (VRRP)

---

<b>Syntax</b>	<code>hold-time seconds;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.
<b>Default</b>	VRRP preemption is not timed.
<b>Options</b>	<b><i>seconds</i></b> —Hold-time period. <b>Range:</b> 0 through 3600 seconds <b>Default:</b> 0 seconds (VRRP preemption is not timed.)
<b>Required Privilege Level</b>	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a Backup Router to Preempt the VRRP Master Router</i></li><li>• <i>Configuring VRRP Preemption and Hold Time</i></li></ul>




## inherit-advertisement-interval

---

<b>Syntax</b>	<code>inherit-advertisement-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.2R3.
<b>Description</b>	Set the time interval for advertisement for inherit sessions.
<b>Options</b>	<b>inherit-advertisement-interval <i>seconds</i></b> —Time interval for inherit sessions advertisements in seconds. The default value is the recommended value. <b>Default:</b> 120 <b>Range:</b> 5 to 120
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>•</li></ul>

## inet6-advertise-interval


---

<b>Syntax</b>	<code>inet6-advertise-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <a href="#">vrrp-inet6-group <i>group-id</i></a>],</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> <a href="#">vrrp-inet6-group <i>group-id</i></a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4R2.
<b>Description</b>	<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> <b>NOTE:</b> 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>	
<b>Options</b>	<p><b><i>milliseconds</i></b>—Interval, in milliseconds, between advertisement packets.</p> <p><b>Range:</b> 100 to 40,000 milliseconds (ms)</p> <p><b>Default:</b> 1 second</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Advertisement Interval for the VRRP Master Router</a></li><li>• <a href="#">advertise-interval on page 168</a></li><li>• <a href="#">fast-interval on page 174</a></li><li>• <a href="#">version-3 on page 190</a></li></ul>

## interface

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     <b>bandwidth-threshold</b> <i>bits-per-second</i> <b>priority-cost</b> <i>priority</i>;     <b>priority-cost</b> <i>priority</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <i>group-id</i> <b>track</b>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> <b>track</b>],</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> <b>vrrp-group</b> <i>group-id</i> <b>track</b>],</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> <b>vrrp-inet6-group</b> <i>group-id</i> <b>track</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>bandwidth-threshold</b> statement added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.
<b>Options</b>	<p><b><i>interface-name</i></b>—Interface to be tracked for this VRRP group.</p> <p><b>Range:</b> 1 through 10 interfaces</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring a Logical Interface to Be Tracked for a VRRP Group</i></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> </ul>

## preempt (VRRP)

<b>Syntax</b>	(preempt   no-preempt) { hold-time seconds; }
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
<b>Description</b>	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"> <li>• <b>preempt</b>—Allow the master router to be preempted.</li> </ul> <p>.....</p> <div>  <p><b>NOTE:</b> By default, a higher-priority backup router can preempt a lower-priority master router.</p> <p>.....</p> </div> <ul style="list-style-type: none"> <li>• <b>no-preempt</b>—Prohibit the preemption of the master router. When <b>no-preempt</b> is configured, the backup router cannot preempt the master router even if the backup router has a higher priority.</li> </ul> <p>The remaining statement is explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Default</b>	By default the <b>preempt</b> statement is enabled, and a higher-priority backup router preempts a lower-priority master router even if the <b>preempt</b> statement is not explicitly configured.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring a Backup Router to Preempt the VRRP Master Router</i></li> <li>• <i>Configuring VRRP Preemption and Hold Time</i></li> </ul>

## priority (Protocols VRRP)


<b>Syntax</b>	<code>priority <i>priority</i>;</code>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 18.1R1 for the SRX Series devices.</p>
<b>Description</b>	Configure a Virtual Router Redundancy Protocol (VRRP) device's priority for becoming the master default device. The device with the highest priority within the group becomes the master.
<b>Options</b>	<p><b>priority</b>—Device's priority for being elected to be the master device in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p><b>Range:</b> 1 through 255</p> <p><b>Default:</b> 100. If two or more devices have the highest priority in the VRRP group, the device with the VRRP interface that has the highest IP address becomes the master, and the others serve as backups.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Basic VRRP Support</i></li> <li>• <i>Understanding VRRP on SRX Series Devices</i></li> <li>• <i>Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces</i></li> </ul>

## priority-cost (VRRP)

---

<b>Syntax</b>	<code>priority-cost <i>priority</i>;</code>
<b>Hierarchy Level</b>	<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 interface interface-name</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track interface interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track interface interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track interface interface-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series. Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.
<b>Description</b>	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.
<b>Options</b>	<b><i>priority</i></b> —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. <b>Range:</b> 1 through 254
<b>Required Privilege Level</b>	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><i>Configuring a Logical Interface to Be Tracked for a VRRP Group</i></li></ul>

## priority-hold-time

<b>Syntax</b>	<code>priority-hold-time seconds;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <i>group-id</i> <b>track</b>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> <b>track</b>],</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> <b>vrrp-group</b> <i>group-id</i> <b>track</b>],</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> <b>vrrp-inet6-group</b> <i>group-id</i> <b>track</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
<b>Description</b>	<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>
	<p> <b>NOTE:</b> 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>
<b>Options</b>	<p><b>seconds</b>—Minimum length of time that must elapse between dynamic priority changes.</p> <p><b>Range:</b> 0 through 3600 seconds</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring a Logical Interface to Be Tracked for a VRRP Group</i></li> </ul>

## route (Interfaces)

<b>Syntax</b>	<code>route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;</code>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for QFX Series.</p> <p>Statement introduced in Junos OS 12.1 for EX Series switches.</p>
<b>Description</b>	Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.
<b>Options</b>	<p><b><i>prefix</i></b>—Route to be tracked for this VRRP group.</p> <p><b><i>priority-cost priority</i></b>—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><b><i>routing-instance instance-name</i></b>—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 <b>default</b>.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring a Route to Be Tracked for a VRRP Group</i></li> </ul>



## 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.



**NOTE:** The `skew-timer-disable` statement is used when there is only one master router and one backup router in the network.

**Default**    By default, the skew timer is enabled for all the VRRP groups.

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

**Related Documentation**    • *Improving the Convergence Time for VRRP*  
    • *Configuring VRRP to Improve Convergence Time*

## startup-silent-period

---

<b>Syntax</b>	<code>startup-silent-period <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
<b>Description</b>	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.
<b>Options</b>	<b><i>seconds</i></b> —Number of seconds for the startup period. <b>Default:</b> 4 seconds <b>Range:</b> 1 through 2000 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Startup Period for VRRP Operations</i></li></ul>

## traceoptions (Protocols VRRP)

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;microsecond-stamp&gt; &lt;size <i>size</i>&gt;     &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i>;     no-remote-trace; } </pre>
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory <b>/var/log</b>.</p>
<b>Default</b>	If you do not include this statement, no VRRP-specific tracing operations are performed.
<b>Options</b>	<p><b>file <i>filename</i></b>—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 <b>/var/log</b>. By default, VRRP tracing output is placed in the file <b>vrrpd</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file.0</i></b>, then <b><i>trace-file.1</i></b>, 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><b>Range:</b> 0 through 4,294,967,296 files</p> <p><b>Default:</b> 3 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. These are the VRRP-specific tracing options:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—All VRRP tracing operations</li> <li>• <b>database</b>—Database changes</li> <li>• <b>general</b>—General events</li> <li>• <b>interfaces</b>—Interface changes</li> <li>• <b>normal</b>—Normal events</li> </ul>

- **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.

<b>Required Privilege Level</b>	trace—To view this statement in the configuration. trace-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Tracing VRRP Operations</i></li></ul>
------------------------------	--

## track (VRRP)

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>priority-hold-time</b> statement added in Junos OS Release 8.1.</p> <p><b>route</b> statement added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
<b>Description</b>	Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.
<b>Options</b>	The remaining statements are described separately.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring a Logical Interface to Be Tracked for a VRRP Group</li> <li>Configuring a Route to Be Tracked for a VRRP Group</li> </ul>

## version-3

---

<b>Syntax</b>	version-3;
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	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.
- 

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Junos OS Support for VRRPv3</i></li></ul>

## virtual-address

<b>Syntax</b>	<code>virtual-address [ <i>addresses</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 18.1R1 for the SRX Series devices.</p>
<b>Description</b>	Configure the addresses of the devices in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.
<b>Options</b>	<b>addresses</b> —Addresses of one or more devices. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master device for the group.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Basic VRRP Support</i></li> <li>• <i>Understanding VRRP on SRX Series Devices</i></li> <li>• <i>Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces</i></li> </ul>


## virtual-inet6-address

---


<b>Syntax</b>	<code>virtual-inet6-address [ <i>addresses</i> ];</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>virtual-inet6-group</b> <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> <b>virtual-inet6-group</b> <i>group-id</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.
<b>Options</b>	<b>addresses</b> —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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Basic VRRP Support</i></li></ul>



## virtual-link-local-address

<b>Syntax</b>	<code>virtual-link-local-address <i>ipv6-address</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>virtual-link-local-address</b> <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> <b>virtual-link-local-address</b> <i>group-id</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	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><b>NOTE:</b> 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>	
<b>Options</b>	<i>ipv6-address</i> —virtual link-local IPv6 address for VRRP for an IPv6 group. <b>Range:</b> 0 through 255
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Basic VRRP Support</i></li> <li><i>Junos OS Support for VRRPv3</i></li> </ul>

## vrrp-group

<b>Syntax</b>	<pre> vrrp-group <i>group-id</i> {   (accept-data   no-accept-data);   advertise-interval <i>seconds</i>;   global-advertisements-threshold <i>number</i>;   authentication-key <i>key</i>;   authentication-type <i>authentication</i>;   fast-interval <i>milliseconds</i>;   (preempt   no-preempt) {     hold-time <i>seconds</i>;   }   priority <i>number</i>;   track {     interface <i>interface-name</i> {       bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;       priority-cost <i>priority</i>;     }     priority-hold-time <i>seconds</i>;     route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;   }   virtual-address [ <i>addresses</i> ];   vrrp-inherit-from <i>vrrp-group</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],  [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group. As of Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the <b>nonstop-routing</b> statement at the [edit routing-options] or [edit logical system <i>logical-system-name</i> routing-options] hierarchy level.</p>
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> The group identifier that you enter must be different from any other group identifiers that you configured for logical units of this same physical interface.</p> </div> </div>	
<b>Options</b>	<p><b>group-id</b>—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 <b>source-address-filter</b> statement. MAC addresses ranging from 00:00:5e:00:53:00 through 00:00:5e:00:53:ff are reserved for VRRP,</p>

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. See [CLI Explorer](#).

<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Basic VRRP Support</i></li><li>• <i>Configuring VRRP</i></li><li>• <i>Example: Configuring VRRP for Load Sharing</i></li><li>• <a href="#">vrrp-inet6-group on page 196</a></li><li>• <a href="#">nonstop-routing on page 138</a></li></ul>
------------------------------	--

## 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.



**NOTE:** The group identifier that you enter must be different from any other group identifiers that you configured for logical units of this same physical interface.

**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. See [CLI Explorer](#).

**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*

## vrrp-inherit-from

---

**Syntax** vrrp-inherit-from {  
     active-group *group-index*;  
     active-interface *active-interface-name*;  
 }

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

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** VRRP group to follow for the vrrp-group or vrrp-inet6-group.

**Options** *group-index*—Identifier for VRRP active group.  
**Range:** 0 through 255  
*active-interface-name*—Interface name of VRRP active group.

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

**Related Documentation**

- *Understanding VRRP*



## CHAPTER 12

# Operational Commands

- `clear vrrp`
- `request chassis ssb master switch`
- `request system software in-service-upgrade`
- `request system software in-service-upgrade (MX Series 3D Universal Edge Routers and EX9200 Switches)`
- `request system software validate in-service-upgrade`
- `show chassis ssb`
- `show nonstop-routing`
- `show pfe ssb`
- `show system switchover`
- `show task replication`
- `show vrrp`
- `show vrrp track`

## clear vrrp

---

<b>Syntax</b>	clear vrrp (all   interface <i>interface-name</i> )
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Set Virtual Router Redundancy Protocol (VRRP) interface statistics to zero.
<b>Options</b>	<b>all</b> —Clear statistics on all interfaces. <b>interface <i>interface-name</i></b> —Clear statistics on the specified interface only.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show vrrp on page 256</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear vrrp all on page 200</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

clear vrrp all

```
user@host> clear vrrp all
```



## request chassis ssb master switch

<b>Syntax</b>	request chassis ssb master switch <no-confirm>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M20 router only) Control which System and Switch Board (SSB) is master.
<b>Options</b>	<b>no-confirm</b> —(Optional) Do not request confirmation for the switch.
<b>Additional Information</b>	<p>By default, the SSB in slot 0 (SSB0) is the master and the SSB in slot 1 (SSB1) is the backup. If you use this command to change the master, and then restart the chassis software for any reason, the master reverts to the default setting. To change the default master SSB, include the <b>ssb</b> statement at the <b>[edit chassis redundancy]</b> hierarchy level in the configuration. For more information, see the <i>Junos OS Administration Library</i>.</p> <p>The configurations on the two SSBs do not have to be the same, and they are not automatically synchronized. If you configure both SSBs as masters, when the chassis software restarts for any reason, the SSB in slot 0 becomes the master and the one in slot 1 becomes the backup.</p> <p>The switchover from the primary SSB to the backup SSB is immediate. The SSB takes several seconds to reinitialize the Flexible PIC Concentrators (FPCs) and restart the PICs. The interior gateway protocol (IGP) and BGP convergence times depend on the specific network environment.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show chassis ssb on page 236</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request chassis ssb master switch on page 201</a> <a href="#">request chassis ssb master switch no-confirm on page 202</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request chassis ssb master switch

```

user@host> request chassis ssb master switch
warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between system switch boards ? [yes,no] (no) yes

Switch initiated, use "show chassis ssb" to verify

```

### request chassis ssb master switch no-confirm

```
user@host> request chassis ssb master switch no-confirm  
Switch initiated, use "show chassis ssb" to verify
```

## request system software in-service-upgrade

<b>Syntax</b>	<code>request system software in-service-upgrade <i>package-name</i></code> <code>&lt;no-old-master-upgrade&gt;</code> <code>&lt;reboot&gt;</code>
<b>Syntax (QFX Series)</b>	<code>request system software in-service-upgrade <i>package-name</i></code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 12.3R2, 13.1R2, and 13.2R1 for TX Matrix Plus routers.</p> <p>Command introduced in Junos OS Release 13.2 for PTX5000 routers.</p> <p>Command introduced in Junos OS Release 13.2 X51-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 15.1X54-D60 for the ACX5000 line of routers.</p>
<b>Description</b>	<p>Perform a unified in-service software upgrade (ISSU). A unified ISSU enables you to upgrade from one Junos OS Release to another with no disruption on the control plane and with minimal disruption of traffic.</p> <p>On QFX5100 switches, enable nonstop active routing (NSR) and nonstop bridging (NSB).</p> <p>On QFX5200 switches, non-stop active routing (NSR) and nonstop bridging (NSB).</p>
<b>Options</b>	<p><b><i>package-name</i></b>—Location from which the software package or bundle is to be installed. For example:</p> <ul style="list-style-type: none"> <li><b><i>/var/tmp/package-name</i></b>— For a software package or bundle that is being installed from a local directory on the router.</li> <li><b><i>protocol://hostname/pathname/package-name</i></b>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <b><i>protocol</i></b> with one of the following: <ul style="list-style-type: none"> <li><b>ftp</b>—File Transfer Protocol</li> <li><b>http</b>—Hypertext Transfer Protocol</li> <li><b>scp</b>—Secure copy (available only for Canada and U.S. version)</li> </ul> </li> </ul> <p><b>no-old-master-upgrade</b>—(Optional) When the <b>no-old-master-upgrade</b> option is included, after the backup Routing Engine is rebooted with the new software package and a switchover occurs to make it the new master Routing Engine, the former master (new backup) Routing Engine will not be upgraded to the new software. In this case, you must manually upgrade the former master (new backup) Routing Engine. If you do not include the <b>no-old-master-upgrade</b> option, the system will automatically upgrade the former master Routing Engine.</p>



**NOTE:** This option is not available on the QFX5100 and QFX5200 switches.

**reboot**—(Optional) When the **reboot** option is included, the former master (new backup) Routing Engine is automatically rebooted after being upgraded to the new software. When the **reboot** option is not included, you must manually reboot the former master (new backup) Routing Engine using the **request system reboot** command.



**NOTE:** This option is not available on the QFX series.

**Additional Information** The following conditions apply to unified ISSUs:

- Unified ISSU is not supported on every platform. For a list of supported platforms, see *Unified ISSU System Requirements*.
- Unsupported PICs are restarted during a unified ISSU on certain routing devices. For information about supported PICs, see the *Junos OS High Availability Library for Routing Devices*.
- Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the *Junos OS High Availability Library for Routing Devices*.
- During a unified ISSU, you cannot bring any PICs online or offline on certain routing devices.

For more information, see the *Junos OS High Availability Library for Routing Devices*.

**Required Privilege Level**

view

**Related Documentation**

- *request system software abort*
- *show chassis in-service-upgrade*
- *Getting Started with Unified In-Service Software Upgrade*
- *Performing an In-Service Software Upgrade (ISSU) with Non-Stop Routing*
- *Understanding In-Service Software Upgrade with Graceful Restart*
- *Performing an In-Service Software Upgrade (ISSU) with Graceful Restart*
- *Example: Performing a Unified ISSU*

**List of Sample Output** [request system software-in-service upgrade reboot on page 205](#)

[request system software-in-service upgrade reboot \(TX Matrix Plus Router\) on page 207](#)  
[request system software-in-service upgrade \(QFX5100 Switch\) on page 215](#)

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

## Sample Output

### request system software-in-service upgrade reboot

```
{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.

```

#### request system software-in-service upgrade reboot (TX Matrix Plus Router)

```

{master}

user@host> request system software in-service-upgrade
/var/tmp/jinstall-12.3R2-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
PIC 8/1 will be offlined (In-Service-Upgrade not supported)
PIC 19/2 will be offlined (In-Service-Upgrade not supported)
PIC 15/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

Checking compatibility with configuration
Initializing...
Using jbase-12.3R2
Verified manifest signed by PackageProduction_12_3_0
Using /var/tmp/jinstall-12.3R2-domestic-signed.tgz
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Using jinstall-12.3R2-domestic.tgz
Using jbundle-12.3R2-domestic.tgz
Checking jbundle requirements on /
Using jbase-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jbase-12.3R2 signed by PackageProduction_12_3_0
Using /var/validate/chroot/tmp/jbundle/jboot-12.3R2.tgz
Using jcrypto-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jcrypto-12.3R2 signed by PackageProduction_12_3_0
Using jdocs-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jdocs-12.3R2 signed by PackageProduction_12_3_0
Using jkernel-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jkernel-12.3R2 signed by PackageProduction_12_3_0
Using jpfe-12.3R2.tgz
WARNING: jpfe-12.3R2.tgz: not a signed package
WARNING: jpfe-common-12.3R2.tgz: not a signed package
Verified jpfe-common-12.3R2 signed by PackageProduction_12_3_0
WARNING: jpfe-T-12.3R2.tgz: not a signed package
Verified jpfe-T-12.3R2 signed by PackageProduction_12_3_0
Using jplatform-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jplatform-12.3R2 signed by PackageProduction_12_3_0
Using jroute-12.3R2.tgz

```

```
Verified manifest signed by PackageProduction_12_3_0
Verified jroute-12.3R2 signed by PackageProduction_12_3_0
Using jruntime-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jruntime-12.3R2 signed by PackageProduction_12_3_0
Using jservices-12.3R2.tgz
Using jservices-crypto-12.3R2.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
ISSU: Preparing LCC Backup REs
Pushing bundle to lcc0-re1
Pushing bundle to lcc1-re1
Pushing bundle to lcc2-re1
Pushing bundle to lcc3-re1
Pushing bundle to sfc0-re1
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...  
Saving state for rollback ...  
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...  
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0  
Adding jinstall...  
Verified manifest signed by PackageProduction\_12\_3\_0

WARNING: This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...  
Saving state for rollback ...  
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...  
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0  
Adding jinstall...  
Verified manifest signed by PackageProduction\_12\_3\_0

WARNING: This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...  
Saving state for rollback ...  
ISSU: Preparing SFC Backup RE

```
NOTICE: Validating configuration against jinstall-12.3R2-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-12.3R2
Verified manifest signed by PackageProduction_12_3_0
Using /var/tmp/jinstall-12.3R2-domestic-signed.tgz
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Using jinstall-12.3R2-domestic.tgz
Using jbundle-12.3R2-domestic.tgz
Checking jbundle requirements on /
Using jbase-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jbase-12.3R2 signed by PackageProduction_12_3_0
Using /var/validate/chroot/tmp/jbundle/jboot-12.3R2.tgz
Using jcrypto-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jcrypto-12.3R2 signed by PackageProduction_12_3_0
Using jdocs-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jdocs-12.3R2 signed by PackageProduction_12_3_0
Using jkernel-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jkernel-12.3R2 signed by PackageProduction_12_3_0
Using jpfe-12.3R2.tgz
WARNING: jpfe-12.3R2.tgz: not a signed package
WARNING: jpfe-common-12.3R2.tgz: not a signed package
Verified jpfe-common-12.3R2 signed by PackageProduction_12_3_0
WARNING: jpfe-T-12.3R2.tgz: not a signed package
Verified jpfe-T-12.3R2 signed by PackageProduction_12_3_0
Using jplatform-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jplatform-12.3R2 signed by PackageProduction_12_3_0
Using jroute-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jroute-12.3R2 signed by PackageProduction_12_3_0
Using jruntime-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jruntime-12.3R2 signed by PackageProduction_12_3_0
Using jservices-12.3R2.tgz
Using jservices-crypto-12.3R2.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...

Saving state for rollback ...

SFC Backup upgrade done

Rebooting SFC Backup RE

Rebooting sfc0-re1

ISSU: SFC Backup RE Prepare Done

Waiting for SFC Backup RE reboot

Rebooting lcc0-re1

Rebooting LCC [lcc0-re1]

Rebooting lcc1-re1

Rebooting LCC [lcc1-re1]

Rebooting lcc2-re1

Rebooting LCC [lcc2-re1]

Rebooting lcc3-re1

Rebooting LCC [lcc3-re1]

LCC Backup REs have rebooted

Waiting for LCC Backup REs come back online

ISSU: LCC Backup REs Prepare Done

GRES operational

Initiating Chassis In-Service-Upgrade

Chassis ISSU Started

ISSU: Preparing Daemons

ISSU: Daemons Ready for ISSU

ISSU: Starting Upgrade for FRUs

ISSU: Preparing for Switchover

ISSU: Ready for Switchover

Checking In-Service-Upgrade status

lcc0-re0:

Item	Status	Reason
FPC 1	Online (ISSU)	
PIC 0	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
PIC 1	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	

lcc1-re0:

Item	Status	Reason
FPC 0	Online (ISSU)	
PIC 3	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 6	Online (ISSU)	

FPC 7            Online (ISSU)

lcc2-re0:

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
PIC 0	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	
PIC 1	Online (ISSU)	

lcc3-re0:

Item	Status	Reason
FPC 0	Online (ISSU)	
PIC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
PIC 2	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 5	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	
PIC 1	Online (ISSU)	

lcc0-re0:

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

lcc1-re0:

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

lcc2-re0:

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

lcc3-re0:

Resolving mastership...  
Complete. The other routing engine becomes the master.  
Resolving mastership...  
Complete. The other routing engine becomes the master.  
ISSU: RE switchover Done  
ISSU: Upgrading SFC Old Master RE

lcc0-re0:

Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...  
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0  
Adding jinstall...  
Verified manifest signed by PackageProduction\_12\_3\_0

WARNING:    This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...

lcc1-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:      This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...

lcc2-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:      This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...

lcc3-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:  This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/paBWTg' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:  This package will load JUNOS 12.3R2 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-12.3R2-domestic-signed ...
cp: /var/tmp/paBWTg is a directory (not copied).
Saving state for rollback ...
ISSU: SFC Old Master Upgrade Done
ISSU: IDLE

```

### request system software-in-service upgrade (QFX5100 Switch)

```

{master}

user@switch> request system software in-service-upgrade
/var/tmp/jinstall-qfx-132_x51_vjunos.0-domestic.tgz
ISSU: Validating Image
Prepare for ISSU
spawn the backup VM
ISSU: Preparing Backup RE
Backup upgrade done
ISSU: Backup RE Prepare Done
waiting for backup RE switchover ready
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
send ISSU done to chassisd on backup VM
Chassis ISSU Completed
ISSU: IDLE
mgd_package_opus_issu: Initiate em0 device handoff

```

## request system software in-service-upgrade (MX Series 3D Universal Edge Routers and EX9200 Switches)

---

**Syntax**    `request system software in-service-upgrade package-name`  
                  `<no-copy>`  
                  `<no-old-master-upgrade>`  
                  `<reboot>`  
                  `<unlink>`

**Release Information**    Command introduced in Junos OS Release 11.2.  
                              Command introduced in Junos OS Release 14.1 for MX Series Virtual Chassis.  
                              Command introduced in Junos OS Release 14.2 for EX Series switches.

**Description**    Perform a unified in-service software upgrade (unified ISSU). Unified ISSU enables you to upgrade from one Junos OS release to another with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is supported only by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

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

- ***/var/tmp/package-name***—For a software package or bundle that is being installed from a local directory on the router.
- ***protocol://hostname/pathname/package-name***—For a software package or bundle that is to be downloaded and installed from a remote location. Replace ***protocol*** with one of the following:
  - **ftp**—File Transfer Protocol
  - **http**—Hypertext Transfer Protocol
  - **scp**—Secure copy (available only for Canada and U.S. version)

**no-copy**—(Optional) When the **no-copy** option is included, copies of package files are not saved on the Packet Forwarding Engine.

The **no-copy** option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

**no-old-master-upgrade**—(Optional) When the **no-old-master-upgrade** option is included, after the backup Routing Engine is rebooted with the new software package and a switchover occurs to make it the new master Routing Engine, the former master (new backup) Routing Engine is not upgraded to the new software. In this case, you must manually upgrade the former master (new backup) Routing Engine. If you do not include the **no-old-master-upgrade** option, the system automatically upgrades the former master Routing Engine.



The **no-old-master-upgrade** option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

**reboot**—(Optional) When the **reboot** option is included, the former master (new backup) Routing Engine is automatically rebooted after being upgraded to the new software. When the **reboot** option is not included, you must manually reboot the former master (new backup) Routing Engine using the **request system reboot** command.

The **reboot** option is accepted but ignored for an MX Series Virtual Chassis or an EX9200 Virtual Chassis. A unified ISSU in an MX Series Virtual Chassis or EX9200 Virtual Chassis always reboots all Routing Engines in the member routers or switches.

**unlink**—(Optional) When the **unlink** option is included, the package is removed from **/var/home** whether the installation is successful or unsuccessful.

The **unlink** option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

**Additional Information** The following conditions apply to unified ISSUs:

- Unified ISSUs are supported on MX Series 3D Universal Edge Routers and EX9200 switches.
- Unsupported PICs (on EX9200, PICs are known as “line cards”) are restarted during a unified ISSU. For information about supported PICs, see the *Junos OS High Availability Library for Routing Devices*. For information about supported EX9200 line cards, see *Unified ISSU System Requirements*.
- Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the *Junos OS High Availability Library for Routing Devices* or, for EX9200, see *Unified ISSU System Requirements*.
- During a unified ISSU, you cannot bring any PICs online or offline.

For more information, see the *Junos OS High Availability Library for Routing Devices*.

**Required Privilege Level** view

**Related Documentation**

- *request system software abort*
- *show chassis in-service-upgrade*

**List of Sample Output** [request system software in-service-upgrade reboot on page 218](#)  
[request system software in-service-upgrade \(MX Series Virtual Chassis\) on page 228](#)

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

## Sample Output

### request system software in-service-upgrade reboot

```
{master}

user@host> request system software in-service-upgrade
/var/tmp/jinstall-11.2B2.1-domestic-signed.tgz reboot
Chassis ISSU Check Done
ISSU: Validating Image
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...
Creating /opt/sdk/service-packages/jservices-voice ...
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg ->
/var/sw/pkg/jservices-voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-bgf ...
Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic ->
```

```

/var/sw/pkg/jservices-bgf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-aac1 ...
Removing /opt/sdk/service-packages/jservices-aac1 ...
Removing jservices-aac1-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aac1 ...
Verified jservices-aac1-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aac1 ...
Storing jservices-aac1-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aac1/jservices-aac1-pic ->
/var/sw/pkg/jservices-aac1-pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic ->
/var/sw/pkg/jservices-llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ptsp ...
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic ->
/var/sw/pkg/jservices-ptsp-pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic ->
/var/sw/pkg/jservices-sfw-pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-nat ...
Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic ->
/var/sw/pkg/jservices-nat-pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...
Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic ->
/var/sw/pkg/jservices-alg-pic-11.2B2.1.tgz...

```

```
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic ->
/var/sw/pkg/jservices-cpcd-pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic ->
/var/sw/pkg/jservices-rpm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic ->
/var/sw/pkg/jservices-hcm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic ->
/var/sw/pkg/jservices-appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-idp ...
Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic ->
/var/sw/pkg/jservices-idp-pic-11.2B2.1.tgz...
Using jservices-crypto-11.2B2.1.tgz
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-crypto-base ...
Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic
-> /var/sw/pkg/jservices-crypto-base-pic-11.2B2.1.tgz...
```

```

Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic ->
/var/sw/pkg/jservices-ssl-pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
ISSU: Preparing Backup RE
Pushing bundle to re1
NOTICE: Validating configuration against jinstall-11.2B2.1-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...
Creating /opt/sdk/service-packages/jservices-voice ...
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg ->
/var/sw/pkg/jservices-voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...

```

```
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-bgf ...
Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic ->
/var/sw/pkg/jservices-bgf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-aac1 ...
Removing /opt/sdk/service-packages/jservices-aac1 ...
Removing jservices-aac1-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aac1 ...
Verified jservices-aac1-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aac1 ...
Storing jservices-aac1-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aac1/jservices-aac1-pic ->
/var/sw/pkg/jservices-aac1-pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic ->
/var/sw/pkg/jservices-llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ptsp ...
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic ->
/var/sw/pkg/jservices-ptsp-pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic ->
/var/sw/pkg/jservices-sfw-pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-nat ...
Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic ->
/var/sw/pkg/jservices-nat-pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
```

```

Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...
Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic ->
/var/sw/pkg/jservices-alg-pic-11.2B2.1.tgz...
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic ->
/var/sw/pkg/jservices-cpcd-pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic ->
/var/sw/pkg/jservices-rpm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic ->
/var/sw/pkg/jservices-hcm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic ->
/var/sw/pkg/jservices-appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-idp ...
Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic ->
/var/sw/pkg/jservices-idp-pic-11.2B2.1.tgz...
Using jservices-crypto-11.2B2.1.tgz
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...

```

```
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-crypto-base ...
Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic
-> /var/sw/pkg/jservices-crypto-base-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic ->
/var/sw/pkg/jservices-ssl-pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-11.2B2.1-domestic-signed.tgz' ...
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Adding jinstall...
Verified manifest signed by PackageProduction_11_2_0

WARNING: This package will load JUNOS 11.2B2.1 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-11.2B2.1-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: 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 1         Online (ISSU)
  FPC 4         Online (ISSU)
  FPC 8         Online (ISSU)
  FPC 10        Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
NOTICE: Validating configuration against jinstall-11.2B2.1-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...
Creating /opt/sdk/service-packages/jservices-voice ...
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg ->
/var/sw/pkg/jservices-voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0

```

```
Creating /opt/sdk/service-packages/jservices-bgf ...
Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic ->
/var/sw/pkg/jservices-bgf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-aac1 ...
Removing /opt/sdk/service-packages/jservices-aac1 ...
Removing jservices-aac1-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aac1 ...
Verified jservices-aac1-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aac1 ...
Storing jservices-aac1-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aac1/jservices-aac1-pic ->
/var/sw/pkg/jservices-aac1-pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic ->
/var/sw/pkg/jservices-llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ptsp ...
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic ->
/var/sw/pkg/jservices-ptsp-pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic ->
/var/sw/pkg/jservices-sfw-pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-nat ...
Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic ->
/var/sw/pkg/jservices-nat-pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...
```

```

Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic ->
/var/sw/pkg/jservices-alg-pic-11.2B2.1.tgz...
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic ->
/var/sw/pkg/jservices-cpcd-pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic ->
/var/sw/pkg/jservices-rpm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic ->
/var/sw/pkg/jservices-hcm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic ->
/var/sw/pkg/jservices-appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-idp ...
Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic ->
/var/sw/pkg/jservices-idp-pic-11.2B2.1.tgz...
Using jservices-crypto-11.2B2.1.tgz
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-crypto-base ...

```

```

Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic
-> /var/sw/pkg/jservices-crypto-base-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic ->
/var/sw/pkg/jservices-ssl-pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-11.2B2.1-domestic-signed.tgz' ...
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Adding jinstall...
Verified manifest signed by PackageProduction_11_2_0

WARNING: This package will load JUNOS 11.2B2.1 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-11.2B2.1-domestic-signed.tgz ...
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
Reboot consistency check bypassed - jinstall 11.2B2.1 will complete installation
upon reboot
[pid 66780]

*** FINAL System shutdown message from user@host> ***
System going down IMMEDIATELY

```

### request system software in-service-upgrade (MX Series Virtual Chassis)

```

{master:member0-re0}

user@host> request system software in-service-upgrade
jinstall-14.1-20140114.2-domestic-signed.tgz
[Jan 30 10:45:32]:ISSU: IDLE

```

```

Beginning in-service-upgrade at Jan 30, 2014; 10:45:34
[Jan 30 10:45:34]:ISSU: Validating Image
Validating VC readiness...
Validating required configuration...
Validating release compatibility...
Validation successful
Initiating chassis in-service-upgrade
[Jan 30 10:46:56]:ISSU: Preparing LCC Backup REs
Copying new release to all RE's
Pushing bundle to member0-re0
Pushing bundle to member1-re0
Pushing bundle to member1-re1
[Jan 30 10:51:11]:ISSU: Preparing Backup RE
Arming new release on all RE's
member0-re0:
-----
Installing package
'/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by
PackageDevelopmentEc_2014
Adding jinstall...

WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

verixec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:    This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 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-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...

member1-re0:
-----
Installing package
'/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by
PackageDevelopmentEc_2014
Adding jinstall...

WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

```

```
verixec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING: This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 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-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...

member1-rel:
-----
Installing package
'/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by
PackageDevelopmentEc_2014
Adding jinstall...

WARNING: The software that is being installed has limited support.
WARNING: Run 'file show /etc/notices/unsupported.txt' for details.

verixec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING: This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 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-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...
Installing package
'/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by
PackageDevelopmentEc_2014
Adding jinstall...

WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

verixec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:    This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 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-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...
[Jan 30 11:03:12]:ISSU: Backup RE Prepare Done
Rebooting standby RE's
Sending Reboot Command to member0-re0
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will
complete installation upon reboot
[pid 2757]
Sending Reboot Command to member1-re1
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will
complete installation upon reboot
[pid 2670]
Waiting for standby RE's to boot
[Jan 30 11:18:26]:ISSU: LCC Backup REs Prepare Done
Waiting for standby RE's to have the correct ISSU state
Waiting for protocol backup to be ready to switch mastership
Switching mastership on the protocol backup chassis to slot 1
Waiting for protocol backup chassis master switch to complete
Globally updating ISSU state
Waiting for protocol backup chassis to become GRES ready
[Jan 30 11:19:18]:ISSU: VC Protocol Backup has Switched
Passing ISSU control to chassisd
Chassis ISSU Started
[Jan 30 11:21:01]:ISSU: Preparing Daemons
[Jan 30 11:22:02]:ISSU: Daemons Ready for ISSU

```

```
[Jan 30 11:22:06]:ISSU: Starting Upgrade for FRUs
[Jan 30 11:25:42]:ISSU: Preparing for Switchover
[Jan 30 11:26:06]:ISSU: Ready for Switchover
[Jan 30 11:26:20]:ISSU: All VC Members Ready for Switchover
Waiting for master chassis to be switch ready
Switching mastership locally
Resolving mastership...
Complete. The other routing engine becomes the master.
Waiting for virtual chassis roles to switch
Globally updating ISSU state to IDLE
[Jan 30 11:26:33]:ISSU: IDLE
Rebooting protocol backup standby RE.
Sending Reboot Command to member1-re0

member1-re0:
-----
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will
complete installation upon reboot
[pid 10462]
Rebooting locally to complete the in service upgrade.
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will
complete installation upon reboot
[pid 13458]

{local:member0-re1}
user@host>
*** FINAL System shutdown message from user@host ***

System going down IMMEDIATELY

Connection closed by foreign host.
```



## request system software validate in-service-upgrade

<b>Syntax</b>	<code>request system software validate in-service-upgrade <i>package-name</i></code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 13.2 for PTX5000 routers.</p> <p>Command introduced in Junos OS Release 14.2 for EX Series switches.</p>
<b>Description</b>	<p>Perform a compatibility check to ensure that the software and hardware components and the configuration on the device support unified ISSU. The <b>request system software validate in-service-upgrade</b> command enables you to detect any compatibility issues before actually issuing the <b>request system software in-service-upgrade</b> command to initiate unified ISSU.</p>
<b>Options</b>	<p><b><i>package-name</i></b>—Location from which the software package or bundle is to be installed. For example:</p> <ul style="list-style-type: none"> <li>• <b><i>/var/tmp/package-name</i></b>—For a software package or bundle that is being installed from a local directory on the router.</li> <li>• <b><i>protocol://hostname/pathname/package-name</i></b>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <b><i>protocol</i></b> with one of the following: <ul style="list-style-type: none"> <li>• <b>ftp</b>—File Transfer Protocol</li> <li>• <b>http</b>—Hypertext Transfer Protocol</li> <li>• <b>scp</b>—Secure copy (available only for Canada and U.S. version)</li> </ul> </li> </ul>
<b>Additional Information</b>	<p>Unified ISSU is not supported on every platform. For a list of supported platforms, see <i>Unified ISSU System Requirements</i>.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">No title found in topic database on page 203</a></li> <li>• <i>request system software abort</i></li> <li>• <i>show chassis in-service-upgrade</i></li> <li>• <i>Getting Started with Unified In-Service Software Upgrade</i></li> <li>• <i>Example: Performing a Unified ISSU</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system software validate in-service-upgrade on page 234</a>
<b>Output Fields</b>	When you enter this command, Junos OS displays the status of your request.

## Sample Output

### request system software validate in-service-upgrade

```
{master}

user@host> request system software validate in-service-upgrade
/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz reboot
Checking compatibility with configuration
Initializing...
Using jbase-9.5-20090127.0
Verified manifest signed by PackageProduction_9_5_0
Using /var/tmp/jinstall-9.6-daily-domestic-signed.tgz
Verified jinstall-9.6-20090706.0-domestic.tgz signed by PackageProduction_9_6_0
Using jinstall-9.6-20090706.0-domestic.tgz
Using jbundle-9.6-20090706.0-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jkernel-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jcrypto-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jpfe-9.6-20090706.0.tgz
Using jdocs-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jroute-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jservices-9.6-20090706.0.tgz
[: /var/validate/chroot/tmp/jservices/packages/jservices-voice-9.6-20090706.0.tgz:
  unexpected operator
Auto-deleting old jservices-voice ...
Removing /opt/sdk/jservices-voice ...
Removing jservices-voice-bsg-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /var/sw/pkg ...
Creating /opt/sdk/jservices-voice ...
Storing jservices-voice-bsg-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-voice/jservices-voice-bsg ->
/var/sw/pkg/jservices-voice-bsg-9.6-20090706.0.tgz...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-bgf ...
Storing jservices-bgf-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-bgf/jservices-bgf-pic ->
/var/sw/pkg/jservices-bgf-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-aacl ...
Removing /opt/sdk/jservices-aacl ...
Removing jservices-aacl-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aacl ...
Verified jservices-aacl-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-aacl ...
Storing jservices-aacl-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-aacl/jservices-aacl-pic ->
/var/sw/pkg/jservices-aacl-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/jservices-llpdf ...
Removing jservices-llpdf-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
```

```
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-llpdf ...
Storing jservices-llpdf-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-llpdf/jservices-llpdf-pic ->
/var/sw/pkg/jservices-llpdf-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/jservices-sfw ...
Removing jservices-sfw-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-sfw ...
Storing jservices-sfw-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-sfw/jservices-sfw-pic ->
/var/sw/pkg/jservices-sfw-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/jservices-appid ...
Removing jservices-appid-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-appid ...
Storing jservices-appid-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-appid/jservices-appid-pic ->
/var/sw/pkg/jservices-appid-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/jservices-idp ...
Removing jservices-idp-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-idp ...
Storing jservices-idp-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-idp/jservices-idp-pic ->
/var/sw/pkg/jservices-idp-pic-9.6-20090706.0.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
PIC 7/0 will be offlined (In-Service-Upgrade not supported)
PIC 7/1 will be offlined (In-Service-Upgrade not supported)
PIC 4/2 will be offlined (In-Service-Upgrade not supported)
PIC 4/3 will be offlined (In-Service-Upgrade not supported)
```

## show chassis ssb

<b>Syntax</b>	<code>show chassis ssb</code> <code>&lt;slot&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M20 routers only) Display status information about the System and Switch Board (SSB).
<b>Options</b>	<p><b>none</b>—Display information about all SSBs.</p> <p><b>slot</b>—(Optional) Display information about the SSB in the specified slot. Replace <b>slot</b> with 0 or 1.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">request chassis ssb master switch on page 201</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show chassis ssb on page 237</a>
<b>Output Fields</b>	Table 5 on page 236 lists the output fields for the <code>show chassis ssb</code> command. Output fields are listed in the approximate order in which they appear.

Table 5: show chassis ssb Output Fields

Field Name	Field Description
<b>Failover</b>	Number of times mastership has changed.
<b>Slot</b>	SSB slot number.
<b>State</b>	<p>Current state of the SSB in this slot. State can be any one of the following:</p> <ul style="list-style-type: none"> <li><b>Master</b>—SSB is online, operating as master.</li> <li><b>Backup</b>—SSB running as backup.</li> <li><b>Empty</b>—No SSB is present.</li> </ul>
<b>Temperature</b>	Temperature of the air passing by the SSB, in degrees Celsius.
<b>CPU utilization</b>	Total percentage of the CPU being used by the SSB's processor.
<b>Interrupt utilization</b>	Of the total CPU being used by the SSB's processor, the percentage being used for interrupts.
<b>Heap utilization</b>	Percentage of heap space being used by the SSB's processor.

*Table 5: show chassis ssb Output Fields (continued)*

Field Name	Field Description
<b>Buffer utilization</b>	Percentage of buffer space being used by the SSB's processor.
<b>DRAM</b>	Total DRAM available to the SSB's processor.
<b>Start time</b>	Time when the SSB started running.
<b>Uptime</b>	How long the SSB has been up and running.

## Sample Output

show chassis ssb

```

user@host> show chassis ssb
SSB status:
  Failover:                0 time
  Slot 0:
    State:                  Master
    Temperature:            33 Centigrade
    CPU utilization:         0 percent
    Interrupt utilization:   0 percent
    Heap utilization:        0 percent
    Buffer utilization:       6 percent
    DRAM:                    64 Mbytes
    Start time:              1999-01-15 22:05:36 UTC
    Uptime:                  21 hours, 21 minutes, 22 seconds
...

```

## show nonstop-routing

<b>Syntax</b>	show nonstop-routing
<b>Release Information</b>	Command introduced in Junos OS Release 13.3.
<b>Description</b>	Display the status of nonstop active routing that includes the automerger statistics and state.
<b>Required Privilege Level</b>	View
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">nonstop-routing on page 138</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show nonstop-routing (MX Series Router) on page 239</a> <a href="#">show nonstop-routing (MX Series Router) on page 240</a>
<b>Output Fields</b>	Table 6 on page 238 describes the output fields for the <b>show nonstop-routing</b> command. Output fields are listed in the approximate order in which they appear.

*Table 6: show nonstop-routing Output Fields*

Field Name	Field Description
Nonstop Routing	State of NSR.
Precision Timers state	<p>State of precision timer feature in the kernel.</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—By default, autokeepalive precision timers are enabled on the kernel after switchover.</li> <li>• <b>Disabled</b>—Autokeepalive precision timers are disabled.</li> <li>• <b>Inactive</b>—Precision timer is inactive if it is disabled.</li> <li>• <b>Ready</b>—Kernel precision timer is ready but is never activated.</li> <li>• <b>InProcess</b>—Kernel precision timer is operational and is generating keepalives on behalf of the RPD after switchover. The / count indicates the number of sessions being serviced against the total sessions.</li> <li>• <b>Completed</b>—Kernel has completed keepalive generation for all the sessions after switchover, and RPD has taken over all of them successfully.</li> <li>• <b>Error</b>—Error while retrieving the precision timer status of the kernel.</li> </ul>
Precision Timers max period	Maximum period, in seconds, after the switchover from standby to master event for which the kernel autogenerates keepalives on behalf of BGP.

Table 6: show nonstop-routing Output Fields (continued)

Field Name	Field Description
Automerge	<p>Status of the automerge.</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—Automerge of sockets by the kernel after switchover is active.</li> <li>• <b>Inactive</b>—Automerge of sockets by the kernel after switchover is inactive.</li> </ul>
Batching	<p>Status of Batching.</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—Automerge of sockets by the kernel after a switchover.</li> <li>• <b>No</b>—Automerge of sockets by the kernel after switchover is inactive.</li> </ul>
Batch count	Number of sockets merged per batch.
Batch count adjust	<p>Speed at which the batch count is adjusted.</p> <ul style="list-style-type: none"> <li>• <b>Slow</b>—Number of sockets merged per batch is incremented additively.</li> <li>• <b>Exp</b>—Number of sockets merged per batch is incremented exponentially.</li> <li>• <b>None</b>—Number of sockets merged per batch remains constant.</li> </ul>
Batch interval	Time interval between batches of automerge operation.
Batch interval adjust	<p>Speed at which the batch interval is adjusted.</p> <ul style="list-style-type: none"> <li>• <b>Exp</b>—Time interval between automerge of batches is increased exponentially.</li> <li>• <b>None</b>—Time interval between automerge of batches is not adjusted.</li> </ul>
Automerge State	<p>State of the automerge</p> <ul style="list-style-type: none"> <li>• <b>Ready</b>—Ready to automerge socket pairs from secondary to primary routing engine</li> <li>• <b>InProgress</b>—Kernel is performing automerge after switchover</li> <li>• <b>Switchover Completed</b>—Sessions merged after switchover</li> </ul>
Sessions Processed	Count of sessions that are automergerd.

## Sample Output

### show nonstop-routing (MX Series Router)

```

user@host show nonstop-routing
Nonstop Routing : Enabled
Precision Timers state: Enabled: Completed - 0/0
Precision Timers max period: 200
Automerge : Active

```

```
Batching: No
Batch count: 200
Batch count adjust: Exponential
Batch interval: 20 msec
Batch interval adjust: None
Automerge State: Ready
Sessions Processed: 0
```

### show nonstop-routing (MX Series Router)

```
user@host> show nonstop-routing
```

```
Nonstop Routing : Enabled
  Automerge : Active
  Batching: Yes
  Batch count: 500
  Batch count adjust: Slow
  Batch interval: 50 msec
  Batch interval adjust: None
  Automerge State: Ready
  Sessions Processed: 0
```



## show pfe ssb

**Syntax**    show pfe ssb

**Release Information**    Command introduced before Junos OS Release 7.4.

**Description**    (M20 routers only) Display Packet Forwarding Engine System and Switch Board (SSB) status and statistics information.

**Options**    This command has no options.

**Required Privilege Level**    admin

**List of Sample Output**    [show pfe ssb on page 243](#)

**Output Fields**    [Table 7 on page 241](#) lists the output fields for the **show pfe ssb** command. Output fields are listed in the approximate order in which they appear.

*Table 7: show pfe ssb Output Fields*

Field Name	Field Description
Uptime (total)	SSB uptime.
Failures	Number of failures .
Pending	Number of pending.
Peer message type receive qualifiers	Information about Peer message type receive qualifiers.
Message Type	Peer message type.
Receive Qualifier	Peer receive qualifier.
TTP	Peer message type TTP.
IFD	Peer message type IFD.
IFL	Peer message type IFL.
Nextthop	Peer message type Nextthop.
COS	Peer message type COS.
Route	Peer message type Route.

Table 7: show pfe ssb Output Fields (continued)

Field Name	Field Description
SW Firewall	Peer message type SW Firewall.
HW Firewall	Peer message type HW Firewall.
PFE Statistics	Peer message type PFE Statistics.
PIC Statistics	Peer message type PIC Statistics.
Sampling	Peer message type Sampling .
Monitoring	Peer message type Monitoring.
ASP	Peer message type ASP.
L2TP	Peer message type L2TP.
Collector	Peer message type Collector.
PIC Configuration	Peer message type PIC Configuration.
Queue Statistics	Peer message type Queue Statistics.
PFE Listener statistics	Information about Packet Forwarding Engine listener statistics: <ul style="list-style-type: none"> <li>• <b>Open</b>—Number of PFE listeners in the “open” state.</li> <li>• <b>Close</b>—Number of PFE listeners in the “close” state.</li> <li>• <b>Sleep</b>—Number of PFE listeners in the “sleep” state.</li> <li>• <b>Wakeup</b>—Number of PFE listeners in the “wakeup” state.</li> <li>• <b>Resync Request</b>—Number of PFE listeners in the “resync request” state.</li> <li>• <b>Resync Done</b>—Number of PFE listeners in the “resync done” state.</li> <li>• <b>Resync Fail</b>—Number of PFE listeners in the “resync fail” state</li> <li>• <b>Resync Time</b>—Number of PFE listeners in the resync time state.</li> </ul>

Table 7: show pfe ssb Output Fields (continued)

Field Name	Field Description
PFE IPC statistics	<p>Information about Packet Forwarding Engine IPC statistics.</p> <ul style="list-style-type: none"> <li>• <b>type</b>—Type of IPC message. <ul style="list-style-type: none"> <li>• <b>Header</b>—IPC message type Header.</li> <li>• <b>Test</b>—IPC message type Test.</li> <li>• <b>Interface</b>—IPC message type Interface.</li> <li>• <b>Chassis</b>—IPC message type Chassis.</li> <li>• <b>Boot</b>—IPC message type Boot</li> <li>• <b>Next-hop</b>—IPC message type Next-hop.</li> <li>• <b>Jtree</b>—IPC message type Jtree.</li> <li>• <b>Cprod</b>—IPC message type Cprod.</li> <li>• <b>Route</b>—IPC message type Route.</li> <li>• <b>Pfe</b>—IPC message type PFE.</li> <li>• <b>Dfw</b>—IPC message type Dfw.</li> <li>• <b>Mastership</b>—IPC message type Mastership.</li> <li>• <b>Sampling</b>—IPC message type Sampling.</li> <li>• <b>GUCP</b>—IPC message type GUCP.</li> <li>• <b>CoS</b>—IPC message type CoS.</li> <li>• <b>GCCP</b>—IPC message type GCCP.</li> <li>• <b>GHCP</b>—IPC message type GHCP.</li> <li>• <b>IRSD</b>—IPC message type IRSD.</li> <li>• <b>Monitoring</b>—IPC message type Monitoring.</li> <li>• <b>RE</b>—IPC message type RE.</li> <li>• <b>PIC</b>—IPC message type PIC.</li> <li>• <b>ASP cfg</b>—IPC message type ASP configuration.</li> <li>• <b>ASP cmd</b>—IPC message type ASP command..</li> <li>• <b>L2TP cfg</b>—IPC message type L2TP configuration.</li> <li>• <b>Collector</b>—IPC message type Collector.</li> <li>• <b>PIC state</b>—IPC message type PIC state.</li> <li>• <b>Aggregator</b>—IPC message type Aggregate.</li> <li>• <b>Empty</b>—IPC message type Empty.</li> <li>• <b>PFE socket-buffer mbuf depth</b>—Information about Packet Forwarding Engine socket-buffer depth</li> <li>• <b>bucket</b>—mbuf bucket value.</li> <li>• <b>count</b>—mbuf count value.</li> </ul> </li> </ul>
PFE socket-buffer bytes pending transmit—	<p>Information about Packet Forwarding Engine socket-buffer bytes pending for transmit.</p> <ul style="list-style-type: none"> <li>• <b>TX Messages</b>—Number of transmitted messages.</li> <li>• <b>RX messages</b>—Number of received messages.</li> </ul>

## Sample Output

### show pfe ssb

```
user@host> show pfe ssb
```

## SSB status:

```

Slot:          Present
State:         Online
Last State Change: 2005-03-06 03:10:28 PST
Uptime (total): 11:23:27
Failures:      0
Pending:       0

```

## Peer message type receive qualifiers:

Message Type	Receive Qualifier
-----	
TTP	Slot only
IFD	All
IFL	All
Nexthop	All
COS	All
Route	All
SW Firewall	All
HW Firewall	All
PFE Statistics	All
PIC Statistics	None
Sampling	All
Monitoring	None
ASP	None
L2TP	None
Collector	None
PIC Configuration	None
Queue Statistics	None
(null)	None

## PFE listener statistics:

```

Open:          1
Close:         0
Sleep:         0
Wakeup:        0
Resync Request: 0
Resync Done:   1
Resync Fail:   0
Resync Time:   0

```

## PFE IPC statistics:

type	TX Messages	RX messages
-----		
Header	0	0
Test	0	0
Interface	737	9911
Chassis	0	0
Boot	0	0
Next-hop	48	0
Jtree	0	0
Cprod	0	0
Route	94	0
Pfe	2034	683
Dfw	8	0
Mastership	0	0
Sampling	0	0
GUCP	0	0
CoS	73	0
GCCP	0	0

GHCP	0	0
IRSD	0	0
Monitoring	0	0
RE	0	0
PIC	0	0
ASP cfg	0	0
ASP cmd	0	0
L2TP cfg	0	0
Collector	0	0
PIC state	0	0
Aggregator	0	0
Empty	0	0

PFE socket-buffer mbuf depth:

bucket	count
0	0
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	0
9	0
10	0
11	0
12	0
13	0
14	0
15	0
16	0
17	0
18	0
19	0
20	0
21	0

PFE socket-buffer bytes pending transmit:

bucket	count
0	0
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	0
9	0
10	0
11	0
12	0
13	0
14	0
15	0
16	0
17	0
18	0

19	0
20	0
21	0

## show system switchover

<b>List of Syntax</b>	<a href="#">Syntax on page 247</a> <a href="#">Syntax (TX Matrix Router) on page 247</a> <a href="#">Syntax (TX Matrix Plus Router) on page 247</a> <a href="#">Syntax (MX Series Router) on page 247</a>
<b>Syntax</b>	show system switchover
<b>Syntax (TX Matrix Router)</b>	show system switchover <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system switchover <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system switchover <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.
<b>Description</b>	Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.



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

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

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

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

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

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

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

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

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

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

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

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

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



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

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

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

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

**Required Privilege Level** view

**Related Documentation**

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

**List of Sample Output**

[show system switchover \(Backup Routing Engine - Ready\) on page 250](#)  
[show system switchover \(Backup Routing Engine - Not Ready\) on page 251](#)  
[show system switchover \(MX Virtual Chassis\) on page 251](#)  
[show system switchover \(MX Virtual Chassis\) on page 251](#)  
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Master Ready on page 251](#)  
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Master Not Ready on page 252](#)  
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Backup Ready on page 252](#)  
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Backup Not Ready on page 252](#)  
[show system switchover all-lcc \(Routing Matrix and Routing Matrix Plus\) on page 253](#)

**Output Fields** [Table 8 on page 249](#) describes the output fields for the **show system switchover** command. Output fields are listed in the approximate order in which they appear.

*Table 8: show system switchover Output Fields*

Field Name	Field Description
Graceful switchover	Display graceful Routing Engine switchover status: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates <b>graceful-switchover</b> is specified for the <b>routing-options</b> configuration command.</li> <li>• <b>Off</b>—Indicates <b>graceful-switchover</b> is not specified for the <b>routing-options</b> configuration command.</li> </ul>

Table 8: show system switchover Output Fields (continued)

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

## Sample Output

### show system switchover (Backup Routing Engine - Ready)

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

Switchover Status: Ready is the way the last line of the output reads if you are running Junos OS Release 16.1R1 or later. If you are running Junos OS Release 15.x, the last line of the output reads as Switchover Ready, for example:

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

**show system switchover (Backup Routing Engine - Not Ready)**

```

user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Not Ready

```

Switchover Status: Not Ready is the way the last line of the output reads if you are running Junos OS Release 16.1R1 or later. If you are running Junos OS Release 15.x, the last line of the output reads as Not ready for mastership switch, try after xxx secs, for example:

```

user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Not ready for mastership switch, try after xxx secs.

```

**show system switchover (MX Virtual Chassis)**

```

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

member1:
-----
Command is not applicable on this member of the virtual-chassis

```

**show system switchover (MX Virtual Chassis)**

```

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

member1:
-----
Command is not applicable on this member of the virtual-chassis

```

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

```

user@host> show system switchover
lcc0-re1:
-----
Multichassis replication: On
Configuration database: Ready

```

```
Kernel database: Ready
Peer state: Steady State
Switchover Status: Ready
```

```
lcc2-re0:
```

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

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

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

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

```
lcc2-re1:
```

```
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Not Ready
```

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

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

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

```
lcc0-re0:
```

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

```
lcc2-re1:
```

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

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

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

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Status: Not Ready
```

```
lcc0-re0:
```

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

```
lcc2-re1:
```

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

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

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

```
lcc0-re0:
```

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

```
lcc2-re0:
```

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

## show task replication

**Syntax** `show task replication`

**Release Information** Command introduced in Junos OS Release 8.5.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.  
 Support for logical systems introduced in Junos OS Release 13.3

**Description** Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.



**CAUTION:** If BGP is configured, before attempting nonstop active routing switchover, check the output of `show bgp replication` to confirm that BGP routing table synchronization has completed on the backup Routing Engine. The complete status in the output of `show task replication` only indicates that the socket replication has completed and the BGP synchronization is in progress.

To determine whether BGP synchronization is complete, you must check the **Protocol state** and **Synchronization state** fields in the output of `show bgp replication` on the master Routing Engine. The **Protocol state** must be **idle** and the **Synchronization state** must be **complete**. If you perform NSR switchover before the BGP synchronization has completed, the BGP session might flap.

**Options** This command has no options.

**Required Privilege Level** view

**List of Sample Output** [show task replication \(Issued on the Master Routing Engine\) on page 255](#)  
[show task replication \(Issued on the Backup Routing Engine\) on page 255](#)

**Output Fields** [Table 9 on page 254](#) lists the output fields for the `show task replication` command. Output fields are listed in the approximate order in which they appear.

*Table 9: show task replication Output Fields*

Field Name	Field Description
Stateful replication	Displays whether or not graceful Routing Engine switchover is configured. The status can be <b>Enabled</b> or <b>Disabled</b> .

*Table 9: show task replication Output Fields (continued)*

Field Name	Field Description
<b>RE mode</b>	Displays the Routing Engine on which the command is issued: <b>Master</b> , <b>Backup</b> , or <b>Not applicable</b> (when the router has only one Routing Engine).
<b>Protocol</b>	Protocols that are supported by nonstop active routing.
<b>Synchronization Status</b>	Nonstop active routing synchronization status for the supported protocols. States are <b>NotStarted</b> , <b>InProgress</b> , and <b>Complete</b> .

## Sample Output

### show task replication (Issued on the Master Routing Engine)

```

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

          Protocol      Synchronization Status
          OSPF           NotStarted
          BGP            Complete
          IS-IS          NotStarted
          LDP            Complete
          PIM            Complete

```

### show task replication (Issued on the Backup Routing Engine)

```

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

```

## show vrrp

---

<b>Syntax</b>	<code>show vrrp</code> <code>&lt;brief   detail   extensive   summary&gt;</code> <code>&lt;interface <i>interface-name</i> &lt;group number&gt;&gt;</code> <code>&lt;logical-system <i>logical-system-name</i> &gt;</code> <code>&lt;nsr&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>nsr</b> option added in Junos OS Release 13.2.
<b>Description</b>	Display status information about Virtual Router Redundancy Protocol (VRRP) groups.
<b>Options</b>	<b>none</b> —(Same as <b>brief</b> ) Display brief status information about all VRRP interfaces.  <b>brief   detail   extensive   summary</b> —(Optional) Display the specified level of output.  <b>interface <i>interface-name</i> &lt;group number&gt;</b> —(Optional) Display information and status about the specified VRRP interface and, optionally, the group number.  <b>logical-system <i>logical-system-name</i></b> —(Optional) Perform this operation on a particular logical system.  <b>nsr</b> —(Optional) Display state replication information when graceful Routing Engine switchover (GRES) with nonstop active routing (NSR) is configured. Use only on the backup Routing Engine.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show vrrp track on page 267</a></li><li>• <a href="#">clear vrrp on page 200</a></li></ul>
<b>List of Sample Output</b>	<a href="#">show vrrp on page 262</a> <a href="#">show vrrp brief on page 262</a> <a href="#">show vrrp detail (IPv6) on page 262</a> <a href="#">show vrrp detail (Route Track) on page 263</a> <a href="#">show vrrp detail (Route Track) on page 263</a> <a href="#">show vrrp extensive on page 263</a> <a href="#">show vrrp interface on page 264</a> <a href="#">show vrrp nsr on page 265</a> <a href="#">show vrrp summary on page 266</a>
<b>Output Fields</b>	<a href="#">Table 10 on page 257</a> lists the output fields for the <b>show vrrp</b> command. Output fields are listed in the approximate order in which they appear



Table 10: show vrrp Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the logical interface.	<b>brief extensive none summary</b>
<b>Interface index</b>	Physical interface index number, which reflects its initialization sequence.	<b>extensive</b>
<b>Groups</b>	Total number of VRRP groups configured on the interface.	<b>extensive</b>
<b>Active</b>	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	<b>extensive</b>
<b>Interface VRRP PDU statistics</b>	Non-errored statistics for the logical interface: <ul style="list-style-type: none"> <li>• <b>Advertisement sent</b>—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted.</li> <li>• <b>Advertisement received</b>—Number of VRRP advertisement PDUs received by the interface.</li> <li>• <b>Packets received</b>—Number of VRRP packets received for VRRP groups on the interface.</li> <li>• <b>No group match received</b>—Number of VRRP packets received for VRRP groups that do not exist on the interface.</li> </ul>	<b>extensive</b>
<b>Interface VRRP PDU error statistics</b>	Errored statistics for the logical interface: <ul style="list-style-type: none"> <li>• <b>Invalid IPAH next type received</b>—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets.</li> <li>• <b>Invalid VRRP ttl value received</b>—Number of packets received whose IP time-to-live (TTL) value is not 255.</li> <li>• <b>Invalid VRRP version received</b>—Number of packets received whose VRRP version is not 2.</li> <li>• <b>Invalid VRRP pdu type received</b>—Number of packets received whose VRRP PDU type is not 1.</li> <li>• <b>Invalid VRRP authentication type received</b>—Number of packets received whose VRRP authentication is not none, simple, or md5.</li> <li>• <b>Invalid VRRP IP count received</b>—Number of packets received whose VRRP IP count exceeds 8.</li> <li>• <b>Invalid VRRP checksum received</b>—Number of packets received whose VRRP checksum does not match the calculated one.</li> </ul>	<b>extensive</b>
<b>Physical interface</b>	Name of the physical interface.	<b>detail extensive</b>
<b>Unit</b>	Logical unit number.	All levels
<b>Address</b>	Address of the physical interface.	<b>brief detail extensive none</b>
<b>Index</b>	Physical interface index number, which reflects its initialization sequence.	<b>detail extensive</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive</b>

Table 10: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
VRRP-Traps	Status of VRRP traps: <b>Enabled</b> or <b>Disabled</b> .	<b>detail extensive</b>
VRRP-Version	VRRP version: <b>2</b> or <b>3</b> .	<b>detail extensive</b>
Type and Address	Identifier for the address and the address itself: <ul style="list-style-type: none"> <li>• <b>lcl</b>—Configured local interface address.</li> <li>• <b>mas</b>—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router.</li> <li>• <b>vip</b>—Configured virtual IP addresses.</li> </ul>	<b>brief none summary</b>
Interface state/Int state/State	State of the physical interface: <ul style="list-style-type: none"> <li>• <b>down</b>—The device is present and the link is unavailable.</li> <li>• <b>not present</b>—The interface is configured, but no physical device is present.</li> <li>• <b>unknown</b>—The VRRP process has not had time to query the kernel about the state of the interface.</li> <li>• <b>up</b>—The device is present and the link is established.</li> </ul>	<b>brief extensive none summary</b>
Group	VRRP group number.	<b>brief extensive none summary</b>
State	The state of the interface on which VRRP is running: <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>init</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>master(ISSU)</b>—The master router interface is going through a unified in-service software upgrade.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul>	<b>extensive</b>
VRRP Mode	If the interface inherits its state and configuration from the active VRRP group, or if it is part of the active VRRP group. <ul style="list-style-type: none"> <li>• <b>Active</b>—Part of the active VRRP group</li> <li>• <b>Inherit</b>—Inherits state and configuration from the active VRRP group.</li> </ul>	<b>detail extensive</b>
Priority	Configured VRRP priority for the interface.	<b>detail extensive</b>
Advertisement interval	Configured VRRP advertisement interval.	<b>detail extensive</b>
Authentication type	Configured VRRP authentication type: <b>none</b> , <b>simple</b> , or <b>md5</b> .	<b>detail extensive</b>

Table 10: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
<b>Advertisement Threshold</b>	A value from 1 through 15, used for setting the time when a peer should be considered down.  <ul style="list-style-type: none"> <li>The time a peer is considered down is equal to the advertisement-threshold multiplied by the advertisement-interval.</li> <li>(advertisement-threshold *advertisement-interval) = Peer down.</li> </ul>	<b>detail extensive</b>
<b>Computed Send Rate</b>	How many protocol data units (PDUs) are generated per second.  Based on the number of instances and the advertisement interval.	<b>detail extensive</b>
<b>Preempt</b>	Whether preemption is allowed on the interface: <b>yes</b> or <b>no</b> .	<b>detail extensive</b>
<b>Accept-data mode</b>	Whether the interface is configured to accept packets destined for the virtual IP address: <b>yes</b> or <b>no</b> .	<b>detail extensive</b>
<b>VIP count</b>	Number of virtual IP addresses that have been configured on the interface.	<b>detail extensive</b>
<b>VIP</b>	List of virtual IP addresses configured on the interface.	<b>detail extensive</b>
<b>Advertisement timer</b>	How long, in seconds, until the advertisement timer expires.	<b>detail extensive</b>
<b>Master router</b>	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is <b>N/A</b> .	<b>detail extensive</b>
<b>Virtual router uptime</b>	How long, in seconds, that the virtual router has been up.	<b>detail extensive</b>
<b>Master router uptime</b>	How long, in seconds, that the master route has been up.	<b>detail extensive</b>
<b>Virtual MAC</b>	MAC address associated with the virtual IP address.	<b>detail extensive</b>
<b>Tracking</b>	Whether tracking is <b>enabled</b> or <b>disabled</b> .	<b>detail extensive</b>
<b>Current priority</b>	Current operational priority for being the VRRP master.	<b>detail extensive</b>
<b>Configured priority</b>	Configured base priority for being the VRRP master.	<b>detail extensive</b>
<b>Priority hold-time</b>	Minimum time interval, in seconds, between successive changes to the current priority. <b>Disabled</b> indicates no minimum interval.	<b>detail extensive</b>
<b>Remaining-time</b>	( <b>track</b> option only) Displays the time remaining in the priority hold-time interval.	<b>detail</b>
<b>Interface tracking</b>	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	<b>detail extensive</b>
<b>Interface/Tracked interface/Track Int</b>	Name of the tracked interface.	<b>detail extensive</b>

Table 10: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
<b>Int state/Interface state/State</b>	Current operational state of the tracked interface: <b>up</b> or <b>down</b> .	<b>detail extensive</b>
<b>Int speed/Speed</b>	Current operational speed, in bits per second, of the tracked interface.	<b>detail extensive</b>
<b>Incurred priority cost</b>	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	<b>detail extensive</b>
<b>Threshold</b>	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred.  An entry of <b>down</b> means that the corresponding priority cost is incurred when the interface is down.	<b>detail extensive</b>
<b>Route tracking</b>	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	<b>detail extensive</b>
<b>Route count</b>	The number of routes being tracked.	<b>detail extensive</b>
<b>Route</b>	The IP address of the route being tracked.	<b>detail extensive</b>
<b>VRF name</b>	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	<b>detail extensive</b>
<b>Route state</b>	The state of the route being tracked: <b>up</b> , <b>down</b> , or <b>unknown</b> .	<b>detail extensive</b>
<b>Priority cost</b>	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	<b>detail extensive</b>
<b>Active</b>	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	<b>detail extensive</b>
<b>Group VRRP PDU statistics</b>	Number of VRRP advertisements sent and received by the group.	<b>extensive</b>

Table 10: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Group VRRP PDU error statistics	<p>Errored statistics for the VRRP group:</p> <ul style="list-style-type: none"> <li>• <b>Bad authentication type received</b>—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be <b>none</b>, <b>simple</b>, or <b>md5</b> and must be the same for all routers in the VRRP group.</li> <li>• <b>Bad password received</b>—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group</li> <li>• <b>Bad MD5 digest received</b>—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router.</li> <li>• <b>Bad advertisement timer received</b>—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group.</li> <li>• <b>Bad VIP count received</b>—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance.</li> <li>• <b>Bad VIPADDR received</b>—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance.</li> </ul>	extensive
Group state transition statistics	<p>State transition statistics for the VRRP group:</p> <ul style="list-style-type: none"> <li>• <b>Idle to master transitions</b>—Number of times that the VRRP instance transitioned from the idle state to the master state.</li> <li>• <b>Idle to backup transitions</b>—Number of times that the VRRP instance transitioned from the idle state to the backup state.</li> <li>• <b>Backup to master transitions</b>—Number of times that the VRRP instance transitioned from the backup state to the master state.</li> <li>• <b>Master to backup transitions</b>—Number of times that the VRRP instance transitioned from the master state to the backup state.</li> </ul>	extensive
VR state	<p>The state of the VRRP:</p> <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting, and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>init</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul> <p><b>NOTE:</b> When show vrrp nsr is used on the backup Routing Engine, it displays the current VRRP state on the master Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use on the master Routing Engine.</p>	brief none summary

Table 10: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
NSR	<p>VRRP nonstop active routing is enabled for the configured VRRP group: <b>yes</b> or <b>no</b>.</p> <p><b>NOTE:</b> A <b>yes</b> value means that the new master Routing Engine will immediately start with the VRRP State value from the original master Routing Engine.</p> <p>A <b>no</b> value means that the VRRP session will:</p> <ul style="list-style-type: none"> <li>• Start afresh.</li> <li>• Go through asilent startup period.</li> <li>• Move to a backup state.</li> <li>• Wait for the D Timer to run out before becoming the master (only if the master has not been configured already).</li> </ul>	<b>brief none</b>
RPD-NSR	The routing options have been set to nonstop active routing: <b>yes</b> or <b>no</b> .	<b>brief none</b>
Timer	<p>VRRP timer information:</p> <ul style="list-style-type: none"> <li>• <b>A</b>—How long, in seconds, until the advertisement timer expires.</li> <li>• <b>D</b>—How long, in seconds, until the Master is Down timer expires.</li> </ul>	<b>brief none</b>

## Sample Output

### show vrrp

```

user@host> show vrrp
Interface      State      Group   VR state  Timer   Type   Address
fe-0/0/0.121   up         1       master    A 1.052  1c1    fec0::12:1:1:1
                                     vip     fe80::12:1:1:99
                                     vip     fec0::12:1:1:99
fe-0/0/2.131   up         1       master    A 0.364  1c1    fec0::13:1:1:1
                                     vip     fe80::13:1:1:99
                                     vip     fec0::13:1:1:99

```

### show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see [show vrrp on page 262](#).

### show vrrp detail (IPv6)

```

user@host> show vrrp detail
Physical interface: fe-0/0/0, Unit: 121, Vlan-id: 212, Address: fec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master, VRRP Mode: Active

```

```

Priority: 200, Advertisement interval: 1, Authentication type: none
Advertisement threshold: 3, Computed send rate: 0
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::12:1:1:99,
fec0::12:1:1:99
Advertisement timer: 1.121s, Master router: fe80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

Physical interface: fe-0/0/2, Unit: 131, Vlan-id: 213, Address: fec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::13:1:1:99,
fec0::13:1:1:99
Advertisement timer: 0.327s, Master router: fe80::13:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

#### show vrrp detail (Route Track)

```

user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 1, Vlan-id: 1, Address: 101.1.1.1/24
Index: 324, SNMP ifIndex: 623, VRRP-Traps: enabled, VRRP-Version: 2
Interface state: up, Group: 1, State: master(ISSU), VRRP Mode: Active
Priority: 200, Advertisement interval: 1, Authentication type: none
Advertisement threshold: 3, Computed send rate: 0
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 101.1.1.3
Advertisement Timer: 0.469s, Master router: 101.1.1.1
Virtual router uptime: 00:02:10, Master router uptime: 00:02:05
Virtual Mac: 00:00:5e:00:01:01
Tracking: disabled

```

#### show vrrp detail (Route Track)

```

user@host> show vrrp detail
Physical interface: ge-1/2/0, Unit: 0, Address: 30.30.30.30/24
Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled, VRRP-Version: 2
Interface state: up, Group: 100, State: master
Priority: 150, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
Advertisement timer: 1.218s, Master router: 30.30.30.30
Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
Virtual MAC: 00:00:5e:00:01:64
Tracking: enabled
  Current priority: 150, Configured priority: 150
  Priority hold-time: disabled
  Interface tracking: disabled
  Route tracking: enabled, Route count: 1
    Route      VRF name      Route state      Priority cost
    192.168.40.0/22    default        up                30

```

#### show vrrp extensive

```

user@host> show vrrp extensive
Interface: ge-2/0/0.0, Interface index :65539, Groups: 1, Active :1
Interface VRRP PDU statistics
  Advertisement sent                :0

```

```

    Advertisement received                :0
    Packets received                      :0
    No group match received                :0
Interface VRRP PDU error statistics
    Invalid IPAH next type received        :0
    Invalid VRRP TTL value received         :0
    Invalid VRRP version received          :0
    Invalid VRRP PDU type received         :0
    Invalid VRRP authentication type received:0
    Invalid VRRP IP count received         :0
    Invalid VRRP checksum received         :0

Physical interface: ge-2/0/0, Unit: 0, Address: 10.10.10.1/24
Index: 65539, SNMP ifIndex: 648, VRRP-Traps: enabled, VRRP-Version: 3
Interface state: up, Group: 1, State: backup, VRRP Mode: Active
Priority: 100, Advertisement interval: 1, Authentication type: none
Advertisement threshold: 3, Computed send rate: 0
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 10.10.10.2
Dead timer: 3.078s, Master priority: 0, Master router: 10.10.10.1
Virtual router uptime: 00:00:04
Tracking: disabled
Group VRRP PDU statistics
    Advertisement sent                    :0
    Advertisement received                 :0
Group VRRP PDU error statistics
    Bad authentication Type received       :0
    Bad password received                  :0
    Bad MD5 digest received                :0
    Bad advertisement timer received        :0
    Bad VIP count received                 :0
    Bad VIPADDR received                   :0
Group state transition statistics
    Idle to master transitions              :0
    Idle to backup transitions              :1
    Backup to master transitions            :0
    Master to backup transitions            :0

```

### show vrrp interface

```

user@host> show vrrp interface ge-0/0/0.1
Interface: ge-0/0/0.1, Interface index :324, Groups: 2, Active :2
Interface VRRP PDU statistics
    Advertisement sent                    :39
    Advertisement received                 :0
    Packets received                      :0
    No group match received                :0
Interface VRRP PDU error statistics
    Invalid IPAH next type received        :0
    Invalid VRRP TTL value received         :0
    Invalid VRRP version received          :0
    Invalid VRRP PDU type received         :0
    Invalid VRRP authentication type received:0
    Invalid VRRP IP count received         :0
    Invalid VRRP checksum received         :0

Physical interface: ge-0/0/0, Unit: 1, Vlan-id: 1, Address: 101.1.1.1/24
Index: 324, SNMP ifIndex: 623, VRRP-Traps: enabled, VRRP-Version: 2
Interface state: up, Group: 1, State: master(ISSU), VRRP Mode: Active
Advertisement threshold: 3, Computed send rate: 0
Priority: 200, Advertisement interval: 1, Authentication type: none
Advertisement threshold: 3, Computed send rate: 0

```



```

Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 101.1.1.3
Advertisement Timer: 0.619s, Master router: 101.1.1.1
Virtual router uptime: 00:00:22, Master router uptime: 00:00:17
Virtual Mac: 00:00:5e:00:01:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent                :20
  Advertisement received             :0
Group VRRP PDU error statistics
  Bad authentication Type received   :0
  Bad password received              :0
  Bad MD5 digest received             :0
  Bad advertisement timer received    :0
  Bad VIP count received              :0
  Bad VIPADDR received               :0
Group state transition statistics
  Idle to master transitions          :0
  Idle to backup transitions          :1
  Backup to master transitions        :1
  Master to backup transitions        :0
Interface: fe-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
Interface VRRP PDU statistics
  Advertisement sent                 :      205
  Advertisement received              :         0
  Packets received                   :         0
  No group match received             :         0
Interface VRRP PDU error statistics
  Invalid IPAH next type received     :         0
  Invalid VRRP TTL value received     :         0
  Invalid VRRP version received       :         0
  Invalid VRRP PDU type received      :         0
  Invalid VRRP authentication type received:         0
  Invalid VRRP IP count received      :         0
  Invalid VRRP checksum received      :         0

```

### show vrrp nsr

This command is similar to **show vrrp**. Here, the **VR state** column displays the current VRRP state on the master Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use on the master Routing Engine.

NSR is yes if VRRP nonstop active routing is enabled for the configured VRRP group.

RPD-NSR is yes if the routing options have been set to nonstop active routing.

```

user@host>show vrrp nsr

```

Interface	State	Group	VR state	VR Mode	Type	NSR	RPD-NSR	Address
ge-1/0/1.0	up	1	master	Active	lcl	yes	yes	10.0.0.1
					vip			10.0.0.3
ge-1/0/1.0	up	2	master	Active	lcl	yes	yes	20.0.0.1
					vip			20.0.0.3
ge-1/0/1.0	up	3	master	Active	lcl	yes	yes	30.0.0.1
					vip			30.0.0.3
ge-1/0/1.0	up	4	master	Active	lcl	yes	yes	40.0.0.1

					vip			40.0.0.3
ge-1/0/1.0	up	5	master	Active	lcl	yes	yes	50.0.0.1
					vip			50.0.0.3
ge-1/0/1.0	up	1	master	Active	lcl	yes	yes	1000::1
fe80::200:5eff:fe00:1					vip			
					vip			1000::3
ge-1/0/1.0	up	2	master	Active	lcl	yes	yes	2000::1
fe80::200:5eff:fe00:2					vip			
					vip			2000::3
ge-1/0/1.0	up	3	master	Active	lcl	yes	yes	3000::1
fe80::200:5eff:fe00:3					vip			
					vip			3000::3
ge-1/0/1.0	up	4	master	Active	lcl	yes	yes	4000::1
fe80::200:5eff:fe00:4					vip			
					vip			4000::3
ge-1/0/1.0	up	5	master	Active	lcl	yes	yes	5000::1
fe80::200:5eff:fe00:5					vip			
					vip			5000::3

show vrrp summary

```
user@host> show vrrp summary
```

Interface	State	Group	VR state	Type	Address
ge-4/2/0.0	up	1	backup	lcl	10.57.0.2
				vip	10.57.0.100

## show vrrp track

<b>Syntax</b>	<pre>show vrrp track &lt;all   interfaces   routes&gt; &lt;detail   summary&gt; &lt;logical-system <i>logical-system-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>all</b> and <b>routes</b> options added in Junos OS Release 17.1.</p>
<b>Description</b>	Display status information about Virtual Router Redundancy Protocol (VRRP) tracked routes and tracked interfaces.
<b>Options</b>	<p><b>none</b>—(Same as <b>summary</b>) Display summarized status information of tracked routes and tracked interfaces.</p> <p><b>all   interfaces   routes</b>—(Optional) These options display the following information:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Output is the same as for the <b>show vrrp track</b> command.</li> <li>• <b>interfaces</b>—Show summary of VRRP tracked interfaces.</li> <li>• <b>routes</b>—Show summary of VRRP tracked routes</li> </ul> <p><b>detail   summary</b>—(Optional) Display detailed or summarized information.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Perform this operation on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring a Logical Interface to Be Tracked for a VRRP Group</i></li> <li>• <i>Configuring a Route to Be Tracked for a VRRP Group</i></li> <li>• <a href="#">show vrrp on page 256</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show vrrp track summary on page 269</a></p> <p><a href="#">show vrrp track detail on page 269</a></p> <p><a href="#">show vrrp track interfaces summary on page 269</a></p> <p><a href="#">show vrrp track interfaces detail on page 269</a></p> <p><a href="#">show vrrp track routes summary on page 270</a></p> <p><a href="#">show vrrp track routes detail on page 270</a></p>
<b>Output Fields</b>	<p><a href="#">Table 11 on page 268</a> lists the output fields for the <b>show vrrp track</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 11: show vrrp track Output Fields

Fields	Description	Level
Tracked interface/Track Int	Name of the tracked interface.	detail or summary
State	Current operational state of the tracked interface: <b>up</b> or <b>down</b> .	detail or summary
Speed	Current operational speed, in bits per second, of the tracked interface.	detail or summary
Incurred priority cost	Operational priority cost incurred resulting from the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority cost.	detail
VRRP Int/Tracking VRRP interface	Name of the VRRP interface.	detail or summary
Group	VRRP group number.	detail or summary
VR state	<p>The state of the VRRP:</p> <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting, and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>init</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul> <p><b>NOTE:</b> When the <b>show vrrp nsr</b> command is used on the backup Routing Engine, it displays the current VRRP state on the master Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use the <b>show vrrp nsr</b> command on the master Routing Engine.</p>	detail or summary
Current priority	Current operational priority for being the VRRP master.	detail or summary
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority cost. <b>Disabled</b> indicates no minimum interval.	detail
Track route	IP address of route.	detail or summary
State	State of route. Possible values are <b>unknown</b> , <b>up</b> , and <b>down</b> .	detail or summary
Cost	Priority cost. When the route state is not <b>up</b> , the cost will be deducted from the configured priority of the VRRP session.	detail or summary
Interface	Name of the logical interface (for example, ge-0/0/1.0) on which the corresponding VRRP session is configured.	detail or summary
Cfg	Configured priority.	detail or summary

Table 11: show vrrp track Output Fields (continued)

Fields	Description	Level
Run	Current (or running) priority cost.	detail or summary

## Sample Output

### show vrrp track summary

```
user@host> show vrrp track summary
```

Track Int	State	Speed	VRRP Int	Group	VR State	Current prio
ge-0/0/2.0	up	1g	ge-0/0/1.0	1	master	80
ge-0/0/8.0	up	1g	ge-0/0/1.0	1	master	80

Track route	State	Cost	Interface	Group	Cfg	Run	VR State
44.44.44.0/24	unknown	10	ge-0/0/1.0	1	100	80	master
55.55.55.0/24	unknown	10	ge-0/0/1.0	1	100	80	master

### show vrrp track detail

```
user@host> show vrrp track detail
```

```
Tracked interface: ge-0/0/2.0
State: up, Speed: 1g
Incurred priority cost: 0
Tracking VRRP interface: ge-0/0/1.0, Group: 1
VR State: master
Current priority: 80, Configured priority: 100
Priority hold-time: disabled
```

```
Tracked interface: ge-0/0/8.0
State: up, Speed: 1g
Incurred priority cost: 0
Tracking VRRP interface: ge-0/0/1.0, Group: 1
VR State: master
Current priority: 80, Configured priority: 100
Priority hold-time: disabled
```

Track route	State	Cost	Interface	Group	Cfg	Run	VR State
44.44.44.0/24	unknown	10	ge-0/0/1.0	1	100	80	master
55.55.55.0/24	unknown	10	ge-0/0/1.0	1	100	80	master

### show vrrp track interfaces summary

```
user@host> show vrrp track interfaces summary
```

Track Int	State	Speed	VRRP Int	Group	VR State	Current prio
ge-0/0/2.0	up	1g	ge-0/0/1.0	1	master	80
ge-0/0/8.0	up	1g	ge-0/0/1.0	1	master	80

### show vrrp track interfaces detail

```
user@host> show vrrp track interfaces detail
```

```
Tracked interface: ge-0/0/2.0
State: up, Speed: 1g
Incurred priority cost: 0
Tracking VRRP interface: ge-0/0/1.0, Group: 1
```

```
VR State: master
Current priority: 80, Configured priority: 100
Priority hold-time: disabled
```

```
Tracked interface: ge-0/0/8.0
State: up, Speed: 1g
Incurred priority cost: 0
Tracking VRRP interface: ge-0/0/1.0, Group: 1
VR State: master
Current priority: 80, Configured priority: 100
Priority hold-time: disabled
```

### **show vrrp track routes summary**

```
user@host> show vrrp track routes summary
```

Track route	State	Cost	Interface	Group	Cfg	Run	VR State
44.44.44.0/24	unknown	10	ge-1/0/0.0	1	100	60	bringup
55.55.55.0/24	unknown	10	ge-1/0/0.0	1	100	60	bringup

### **show vrrp track routes detail**

The output for **show vrrp track routes detail** is the same as that for **show vrrp track routes summary**.