



JUNOS® Software

High Availability Configuration Guide

Release 10.0

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Published: 2009-10-15

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

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

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

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

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

Juniper Networks, the Juniper Networks logo, JUNOS, NetScreen, ScreenOS, and Steel-Belted Radius are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

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

JUNOS® Software High Availability Configuration Guide

Release 9.6

Copyright © 2009, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing: Abhilash Prabhakaran

Editing: Nancy Kurahashi

Illustration: Nathaniel Woodward

Cover Design: Edmonds Design

Revision History

October 2009—JUNOS 10.0 R1

The information in this document is current as of the date listed in the revision history.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. The JUNOS Software 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

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

- a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.
- b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.
- c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.
- d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.
- e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE, EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and a copy of the LGPL at <http://www.gnu.org/licenses/lgpl.html>.

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous

agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Abbreviated Table of Contents

About This Guide

xxiii

Part 1	Overview	
Chapter 1	High Availability Overview	3
Part 2	Routing Engine and Switching Control Board Redundancy	
Chapter 2	Routing Engine and Switching Control Board Redundancy Overview	11
Chapter 3	Routing Engine and Switching Control Board Redundancy Configuration Guidelines	21
Chapter 4	Summary of Routing Engine and Switching Control Board Redundancy Statements	33
Part 3	Graceful Routing Engine Switchover	
Chapter 5	Graceful Routing Engine Switchover Overview	45
Chapter 6	Graceful Routing Engine Switchover Configuration Guidelines	53
Chapter 7	Summary of Graceful Routing Engine Switchover Configuration Statements	57
Part 4	Nonstop Bridging	
Chapter 8	Nonstop Bridging Overview	61
Chapter 9	Nonstop Bridging Configuration Guidelines	65
Chapter 10	Summary of Nonstop Bridging Statements	67
Part 5	Nonstop Active Routing	
Chapter 11	Nonstop Active Routing Overview	71
Chapter 12	Nonstop Active Routing Configuration Guidelines	81
Chapter 13	Summary of Nonstop Active Routing Configuration Statements	87
Part 6	Graceful Restart	
Chapter 14	Graceful Restart Overview	95
Chapter 15	Graceful Restart Configuration Guidelines	103
Chapter 16	Summary of Graceful Restart Configuration Statements	141

Part 7	Virtual Router Redundancy Protocol	
Chapter 17	VRRP Overview	155
Chapter 18	VRRP Configuration Guidelines	157
Chapter 19	Summary of VRRP Configuration Statements	175
Part 8	Unified ISSU	
Chapter 20	Unified ISSU Overview	199
Chapter 21	Unified ISSU Configuration Guidelines	217
Chapter 22	Unified ISSU Configuration Statements Summary	235
Part 9	Index	
	Index	241
	Index of Statements and Commands	247

Table of Contents

About This Guide xxiii

JUNOS Documentation and Release Notes	xxiii
Objectives	xxiv
Audience	xxiv
Supported Platforms	xxv
Using the Indexes	xxv
Using the Examples in This Manual	xxv
Merging a Full Example	xxv
Merging a Snippet	xxvi
Documentation Conventions	xxvii
Documentation Feedback	xxviii
Requesting Technical Support	xxix
Self-Help Online Tools and Resources	xxix
Opening a Case with JTAC	xxix

Part 1

Overview

Chapter 1

High Availability Overview 3

Introducing High Availability Features on Juniper Networks Routing	
Platforms	3
Routing Engine Redundancy	3
Graceful Routing Engine Switchover	3
Nonstop Bridging	4
Nonstop Active Routing	4
Graceful Restart	5
Nonstop Active Routing Versus Graceful Restart	6
Effects of a Routing Engine Switchover	6
VRRP	6
Unified ISSU	7
High Availability-Related Features in JUNOS Software	7

Part 2	Routing Engine and Switching Control Board Redundancy	
Chapter 2	Routing Engine and Switching Control Board Redundancy Overview	11
	Understanding Routing Engine Redundancy on JUNOS Routers	11
	Routing Engine Redundancy Overview	11
	Conditions That Trigger a Routing Engine Failover	12
	Default Routing Engine Redundancy Behavior	13
	Routing Engine Redundancy on a TX Matrix Router	14
	Situations That Require You to Halt Routing Engines	15
	Switching Control Board Redundancy	15
	Redundant CFEBs on the M10i Router	16
	Redundant FEBs on the M120 Router	16
	Redundant SSBs on the M20 Router	18
	Redundant SFMs on the M40e and M160 Routers	19
Chapter 3	Routing Engine and Switching Control Board Redundancy Configuration Guidelines	21
	Chassis Redundancy Hierarchy	21
	Initial Routing Engine Configuration Example	22
	Copying a Configuration File from One Routing Engine to the Other	23
	Loading a Software Package from the Other Routing Engine	24
	Configuring Routing Engine Redundancy	24
	Modifying the Default Routing Engine Mastership	25
	Configuring Automatic Failover to the Backup Routing Engine	25
	Without Interruption to Packet Forwarding	25
	On Detection of a Hard Disk Error on the Master Routing Engine	26
	On Detection of a Loss of Keepalive Signal from the Master Routing Engine	26
	When a Software Process Fails	27
	Manually Switching Routing Engine Mastership	27
	Verifying Routing Engine Redundancy Status	27
	Configuring CFEB Redundancy on the M10i Router	28
	Configuring FEB Redundancy on the M120 Router	29
	Example: Configuring FEB Redundancy	30
	Configuring SFM Redundancy on M40e and M160 Routers	31
	Configuring SSB Redundancy on the M20 Router	31
Chapter 4	Summary of Routing Engine and Switching Control Board Redundancy Statements	33
	cfb	33
	description	34
	failover on-disk-failure	34
	failover on-loss-of-keepalives	35

failover other-routing-engine	36
feb	37
feb (Creating a Redundancy Group)	37
feb (Assigning a FEB to a Redundancy Group)	37
keepalive-time	38
no-auto-failover	39
redundancy	39
redundancy-group	40
routing-engine	40
sfm	41
ssb	41

Part 3

Graceful Routing Engine Switchover

Chapter 5

Graceful Routing Engine Switchover Overview **45**

Understanding Graceful Routing Engine Switchover in the JUNOS	
Software	45
Graceful Routing Engine Switchover Concepts	45
Effects of a Routing Engine Switchover	48
Graceful Routing Engine Switchover System Requirements	48
Graceful Routing Engine Switchover Platform Support	49
Graceful Routing Engine Switchover Feature Support	49
Graceful Routing Engine Switchover DPC Support	50
Graceful Routing Engine Switchover and Subscriber Access	51
Graceful Routing Engine Switchover PIC Support	51

Chapter 6

Graceful Routing Engine Switchover Configuration Guidelines **53**

Configuring Graceful Routing Engine Switchover	53
Enabling Graceful Routing Engine Switchover	53
Synchronizing the Routing Engine Configuration	54
Verifying Graceful Routing Engine Switchover Operation	54
Requirements for Routers with a Backup Router Configuration	54
Resetting Local Statistics	55

Chapter 7

Summary of Graceful Routing Engine Switchover Configuration Statements **57**

graceful-switchover	57
---------------------------	----

Part 4	Nonstop Bridging	
Chapter 8	Nonstop Bridging Overview	61
	Nonstop Bridging Concepts	61
	Nonstop Bridging System Requirements	63
	Platform Support	63
	Protocol Support	64
Chapter 9	Nonstop Bridging Configuration Guidelines	65
	Configuring Nonstop Bridging	65
	Enabling Nonstop Bridging	65
	Synchronizing the Routing Engine Configuration	65
	Verifying Nonstop Bridging Operation	66
Chapter 10	Summary of Nonstop Bridging Statements	67
	nonstop-bridging	67
Part 5	Nonstop Active Routing	
Chapter 11	Nonstop Active Routing Overview	71
	Nonstop Active Routing Concepts	71
	Nonstop Active Routing System Requirements	73
	Nonstop Active Routing Platform Support	74
	Nonstop Active Routing Protocol and Feature Support	74
	Nonstop Active Routing BFD Support	76
	Nonstop Active Routing BGP Support	76
	Nonstop Active Routing Layer 2 Circuit and LDP-Based VPLS Support	77
	Nonstop Active Routing PIM Support	77
	Nonstop Active Routing Support for RSVP-TE LSPs	79
Chapter 12	Nonstop Active Routing Configuration Guidelines	81
	Configuring Nonstop Active Routing	81
	Enabling Nonstop Active Routing	81
	Synchronizing the Routing Engine Configuration	82
	Verifying Nonstop Active Routing Operation	82
	Tracing Nonstop Active Routing Synchronization Events	83
	Resetting Local Statistics	84
	Example: Configuring Nonstop Active Routing	84

Chapter 13	Summary of Nonstop Active Routing Configuration Statements	87
	commit synchronize	88
	nonstop-routing	89
	traceoptions	90
 Part 6	 Graceful Restart	
 Chapter 14	 Graceful Restart Overview	 95
	Graceful Restart Concepts	95
	Graceful Restart System Requirements	96
	Aggregate and Static Routes	96
	Graceful Restart and Routing Protocols	96
	BGP	97
	ES-IS	97
	IS-IS	97
	OSPF and OSPFv3	98
	PIM Sparse Mode	98
	RIP and RIPng	98
	Graceful Restart and MPLS-Related Protocols	98
	LDP	99
	RSVP	99
	CCC and TCC	99
	Graceful Restart and Layer 2 and Layer 3 VPNs	100
	Graceful Restart on Logical Systems	101
 Chapter 15	 Graceful Restart Configuration Guidelines	 103
	Configuring Graceful Restart for Aggregate and Static Routes	103
	Configuring Routing Protocols Graceful Restart	103
	Configuring Graceful Restart Globally	104
	Configuring Graceful Restart Options for BGP	104
	Configuring Graceful Restart Options for ES-IS	105
	Configuring Graceful Restart Options for IS-IS	105
	Configuring Graceful Restart Options for OSPF and OSPFv3	106
	Configuring Graceful Restart Options for RIP and RIPng	107
	Configuring Graceful Restart Options for PIM Sparse Mode	107
	Tracking Graceful Restart Events	108
	Configuring Graceful Restart for MPLS-Related Protocols	108
	Configuring Graceful Restart Globally	109
	Configuring Graceful Restart Options for RSVP, CCC, and TCC	109
	Configuring Graceful Restart Options for LDP	110
	Configuring VPN Graceful Restart	111
	Configuring Graceful Restart Globally	111
	Configuring Graceful Restart for the Routing Instance	111

Configuring Logical System Graceful Restart	112
Configuring Graceful Restart Globally	112
Configuring Graceful Restart for a Routing Instance	112
Verifying Graceful Restart Operation	113
Graceful Restart Operational Mode Commands	113
Verifying BGP Graceful Restart	113
Verifying IS-IS and OSPF Graceful Restart	114
Verifying CCC and TCC Graceful Restart	114
Example: Configuring Graceful Restart	115

Chapter 16	Summary of Graceful Restart Configuration Statements	141
-------------------	---	------------

disable	141
graceful-restart	142
helper-disable	143
maximum-helper-recovery-time	143
maximum-helper-restart-time	144
maximum-neighbor-reconnect-time	144
maximum-neighbor-recovery-time	145
no-strict-lsa-checking	145
notify-duration	146
reconnect-time	146
recovery-time	147
restart-duration	148
restart-time	149
stale-routes-time	150
traceoptions	151

Part 7	Virtual Router Redundancy Protocol
---------------	---

Chapter 17	VRRP Overview	155
-------------------	----------------------	------------

Understanding VRRP	155
--------------------------	-----

Chapter 18	VRRP Configuration Guidelines	157
-------------------	--------------------------------------	------------

VRRP Configuration Hierarchy	157
VRRP for IPv6 Configuration Hierarchy	158
Configuring the Startup Period for VRRP Operations	159
Configuring Basic VRRP Support	159
Configuring VRRP Authentication (IPv4 Only)	161
Configuring the Advertisement Interval for the VRRP Master Router	162
Modifying the Advertisement Interval in Seconds	162
Modifying the Advertisement Interval in Milliseconds	163
Configuring a Backup Router to Preempt the Master Router	163
Modifying the Preemption Hold-Time Value	164

Configuring an Interface to Accept Packets Destined for the Virtual IP Address	164
Configuring a Logical Interface to Be Tracked	165
Configuring a Route to Be Tracked	167
Configuring Inheritance for a VRRP Group	167
Tracing VRRP Operations	168
Configuring the Silent Period	169
Configuring Passive ARP Learning for Backup VRRP Routers	169
Enabling the Distributed Periodic Packet Management Process for VRRP	170
Example: Configuring VRRP	171
Example: Configuring VRRP for IPv6	172
Example: Configuring VRRP Route Tracking	173

Chapter 19

Summary of VRRP Configuration Statements **175**

accept-data	175
advertise-interval	176
authentication-key	177
authentication-type	178
bandwidth-threshold	179
fast-interval	180
hold-time	181
inet6-advertise-interval	182
interface	183
no-accept-data	183
no-preempt	183
preempt	184
priority	185
priority-cost	186
priority-hold-time	187
route	188
startup-silent-period	188
traceoptions	189
track	191
virtual-address	192
virtual-inet6-address	192
virtual-link-local-address	193
vrrp-group	194
vrrp-inet6-group	195

Part 8**Unified ISSU****Chapter 20****Unified ISSU Overview****199**

Unified ISSU Concepts	199
Unified ISSU Process on the TX Matrix Router	204
Unified ISSU System Requirements	205
Unified ISSU JUNOS Software Release Support	206
Unified ISSU Platform Support	206
Unified ISSU Protocol Support	206
Unified ISSU Support for Layer 2 Control Protocol Process	207
Unified ISSU Feature Support	208
Unified ISSU PIC Support	208
PIC Considerations	209
SONET/SDH PICs	209
Fast Ethernet and Gigabit Ethernet PICs	211
Channelized PICs	212
Tunnel Services PICs	213
ATM PICs	213
Serial PICs	214
DS3, E1, E3, and T1 PICs	214
Enhanced IQ PICs	215
Enhanced IQ2 Ethernet Services Engine (ESE) PIC	215
Unified ISSU DPC and FPC Support on MX Series Routers	215

Chapter 21**Unified ISSU Configuration Guidelines****217**

Best Practices	217
Before You Begin	217
Verify That the Master and Backup Routing Engines Are Running the Same Software Version	218
Back Up the Router Software	218
Verify That Graceful Routing Engine Switchover and Nonstop Active Routing Are Configured	219
Performing a Unified ISSU	220
Upgrading and Rebooting Both Routing Engines Automatically	220
Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually	225
Upgrading and Rebooting Only One Routing Engine	230
Verifying a Unified ISSU	233
Troubleshooting Unified ISSU Problems	233
Managing and Tracing BFD Sessions During Unified ISSU Procedures	234

Chapter 22**Unified ISSU Configuration Statements Summary****235**

no-issu-timer-negotiation	235
traceoptions	236

Part 9**Index**

Index	241
Index of Statements and Commands	247

List of Figures

Part 3	Graceful Routing Engine Switchover	
Chapter 5	Graceful Routing Engine Switchover Overview	45
	Figure 1: Preparing for a Graceful Routing Engine Switchover	46
	Figure 2: Graceful Routing Engine Switchover Process	47
Part 4	Nonstop Bridging	
Chapter 8	Nonstop Bridging Overview	61
	Figure 3: Nonstop Bridging Switchover Preparation Process	62
	Figure 4: Nonstop Bridging During a Switchover	63
Part 5	Nonstop Active Routing	
Chapter 11	Nonstop Active Routing Overview	71
	Figure 5: Nonstop Active Routing Switchover Preparation Process	72
	Figure 6: Nonstop Active Routing During a Switchover	73
Part 6	Graceful Restart	
Chapter 15	Graceful Restart Configuration Guidelines	103
	Figure 7: Layer 3 VPN Graceful Restart Topology	115
Part 7	Virtual Router Redundancy Protocol	
Chapter 17	VRRP Overview	155
	Figure 8: Basic VRRP	156

List of Tables

	About This Guide	xxiii
	Table 1: Notice Icons	xxvii
	Table 2: Text and Syntax Conventions	xxvii
Part 2	Routing Engine and Switching Control Board Redundancy	
Chapter 3	Routing Engine and Switching Control Board Redundancy Configuration Guidelines	21
	Table 3: Routing Engine Mastership Log	27
Part 3	Graceful Routing Engine Switchover	
Chapter 5	Graceful Routing Engine Switchover Overview	45
	Table 4: Effects of a Routing Engine Switchover	48
	Table 5: Graceful Routing Engine Switchover Feature Support	49
Part 5	Nonstop Active Routing	
Chapter 11	Nonstop Active Routing Overview	71
	Table 6: Nonstop Active Routing Platform Support	74
	Table 7: Nonstop Active Routing Protocol and Feature Support	74
Part 7	Virtual Router Redundancy Protocol	
Chapter 18	VRRP Configuration Guidelines	157
	Table 8: Interface State and Priority Cost Usage	166
Part 8	Unified ISSU	
Chapter 20	Unified ISSU Overview	199
	Table 9: Unified ISSU Platform Support	206
	Table 10: Unified ISSU Protocol Support	207
	Table 11: Unified ISSU PIC Support: SONET/SDH	210
	Table 12: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet	211
	Table 13: Unified ISSU PIC Support: Channelized	212
	Table 14: Unified ISSU PIC Support: Tunnel Services	213
	Table 15: Unified ISSU PIC Support: ATM	213
	Table 16: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC	215

About This Guide

This preface provides the following guidelines for using the *JUNOS® Software High Availability Configuration Guide*:

- JUNOS Documentation and Release Notes on page xxiii
- Objectives on page xxiv
- Audience on page xxiv
- Supported Platforms on page xxv
- Using the Indexes on page xxv
- Using the Examples in This Manual on page xxv
- Documentation Conventions on page xxvii
- Documentation Feedback on page xxviii
- Requesting Technical Support on page xxix

JUNOS Documentation and Release Notes

For a list of related JUNOS documentation, see <http://www.juniper.net/techpubs/software/junos/>.

If the information in the latest release notes differs from the information in the documentation, follow the *JUNOS Software Release Notes*.

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

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using JUNOS Software and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using JUNOS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at <http://www.juniper.net/books>.

Objectives

This guide is designed to provide an overview of high availability concepts and techniques. By understanding the redundancy features of Juniper Networks routing platforms and JUNOS Software, a network administrator can enhance the reliability of a network and deliver highly available services to customers.



NOTE: For additional information about JUNOS Software—either corrections to or information that might have been omitted from this guide—see the software release notes at <http://www.juniper.net/>.

Audience

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M Series, MX Series, T Series, EX Series, or J Series router or switch.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)
- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

Personnel operating the equipment must be trained and competent; must not conduct themselves in a careless, willfully negligent, or hostile manner; and must abide by the instructions provided by the documentation.

Supported Platforms

For the features described in this manual, JUNOS Software currently supports the following platforms:

- J Series
- M Series
- MX Series
- T Series
- EX Series

Using the Indexes

This reference contains two indexes: a complete index that includes topic entries, and an index of statements and commands only.

In the index of statements and commands, an entry refers to a statement summary section only. In the complete index, the entry for a configuration statement or command contains at least two parts:

- The primary entry refers to the statement summary section.
- The secondary entry, *usage guidelines*, refers to the section in a configuration guidelines chapter that describes how to use the statement or command.

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

```

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

```

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

```

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

```

Merging a Snippet

To merge a snippet, follow these steps:

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

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

```

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

```

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

```

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

```

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

```

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

```

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

Documentation Conventions

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

Table 1: Notice Icons





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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the <code>configure</code> command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>JUNOS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>

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

Convention	Description	Examples
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">■ To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.■ The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop address; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">■ In the Logical Interfaces box, select All Interfaces.■ To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select Protocols > Ospf .

Documentation Feedback

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

- Document or topic name
- URL or page number

- Software release version (if applicable)

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <http://www.juniper.net/customers/support/downloads/710059.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

Part 1

Overview

- High Availability Overview on page 3

Chapter 1

High Availability Overview

This chapter contains the following topics:

- Introducing High Availability Features on Juniper Networks Routing Platforms on page 3
- High Availability-Related Features in JUNOS Software on page 7

Introducing High Availability Features on Juniper Networks Routing Platforms

For Juniper Networks routing platforms running JUNOS Software, *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 3
- Graceful Routing Engine Switchover on page 3
- Nonstop Bridging on page 4
- Nonstop Active Routing on page 4
- Graceful Restart on page 5
- Nonstop Active Routing Versus Graceful Restart on page 6
- Effects of a Routing Engine Switchover on page 6
- VRRP on page 6
- Unified ISSU on page 7

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. 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 “Graceful Routing Engine Switchover Overview” on page 45.

Nonstop Bridging

Nonstop bridging enables a routing platform 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 Overview” on page 61.

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 Table 7 on page 74.

For more information on nonstop active routing, see “Nonstop Active Routing Overview” on page 71.

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 either the global **[edit routing-options]** hierarchy level or the hierarchy level for a specific protocol. 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 Overview” on page 95.

Nonstop Active Routing Versus Graceful Restart

Nonstop active routing and graceful restart are two different methods of maintaining high availability. Graceful restart requires a restart process. 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.

In contrast, nonstop active routing does not involve a router restart. Both the master and standby 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.

Effects of a Routing Engine Switchover

Table 4 on page 48 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

For Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, and logical interfaces, you can configure the Virtual Router Redundancy Protocol (VRRP) or VRRP

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

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

For more information, see “Understanding VRRP” on page 155.

Unified ISSU

A unified in-service software upgrade (unified ISSU) enables you to upgrade between two different JUNOS Software 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 “Unified ISSU Overview” on page 199.

Related Topics ■ High Availability-Related Features in JUNOS on page 7

High Availability-Related Features in JUNOS Software

Related redundancy and reliability features include:

- Redundant power supplies, host modules, host subsystems, and forwarding boards. For more information, see the *JUNOS System Basics Configuration Guide* and the *JUNOS 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 Network Interfaces Configuration Guide*.
- Bidirectional Forwarding Detection (BFD) works with other routing protocols to detect failures rapidly. For more information, see the *JUNOS Routing Protocols Configuration Guide*.
- 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 *JUNOS MPLS Applications Configuration Guide*.

- Related Topics** ■ Introducing High Availability Features on Juniper Networks Routing Platforms on page 3

Part 2

Routing Engine and Switching Control Board Redundancy

- Routing Engine and Switching Control Board Redundancy Overview on page 11
- Routing Engine and Switching Control Board Redundancy Configuration Guidelines on page 21
- Summary of Routing Engine and Switching Control Board Redundancy Statements on page 33

Chapter 2

Routing Engine and Switching Control Board Redundancy Overview

For routers that have redundant Routing Engines or redundant switching control boards, including Switching and Forwarding Modules (SFMs), System and Switch Boards (SSBs), Forwarding Engine Boards (FEBs), or Compact Forwarding Engine Board (CFEBs), you can configure redundancy properties. This chapter includes the following topics:

- Understanding Routing Engine Redundancy on JUNOS Routers on page 11
- Switching Control Board Redundancy on page 15

Understanding Routing Engine Redundancy on JUNOS Routers

This topic contains the following sections:

- Routing Engine Redundancy Overview on page 11
- Conditions That Trigger a Routing Engine Failover on page 12
- Default Routing Engine Redundancy Behavior on page 13
- Routing Engine Redundancy on a TX Matrix Router on page 14
- Situations That Require You to Halt Routing Engines on page 15

Routing Engine Redundancy Overview

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.

With redundant Routing Engines, one functions as the master, while the other stands by as a backup should the master Routing Engine fail. When a Routing Engine is configured as master, it has full functionality. It receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding Engine components, and has full control over the chassis. When a Routing Engine is configured to be the backup, it does not communicate with the Packet Forwarding Engine or chassis components.



NOTE: With the introduction of JUNOS Release 8.4, both Routing Engines cannot be configured to be master at the same time. This configuration causes the commit check to fail.

A failover from the master Routing Engine to the backup Routing Engine occurs automatically when the master Routing Engine 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 Routing Engine mastership by issuing one of the **request chassis routing-engine** commands. In this chapter, the term *failover* refers to an automatic event, whereas *switchover* refers to either an automatic or a manual event.

When a failover or a switchover occurs, the backup Routing Engine takes control of the system as the new master Routing Engine:

- If graceful Routing Engine switchover is not configured, when the backup Routing Engine becomes master, it resets the switch plane and downloads its own version of the microkernel to the Packet Forwarding Engine components. Traffic is interrupted while the Packet Forwarding Engine is reinitialized. All kernel and forwarding processes are restarted.
- If graceful Routing Engine switchover is configured, interface and kernel information is preserved. The switchover is faster because the Packet Forwarding Engines are not restarted. The new master Routing Engine restarts the routing protocol process (rpd). All hardware and interfaces are acquired by a process that is similar to a warm restart. For more information about graceful Routing Engine switchover, see “Graceful Routing Engine Switchover Overview” on page 45.
- If graceful Routing Engine switchover and nonstop active routing (NSR) are configured, traffic is not interrupted during the switchover. Interface, kernel, and routing protocol information is preserved. For more information about nonstop active routing, see “Nonstop Active Routing Overview” on page 71.
- If graceful Routing Engine switchover and graceful restart are configured, traffic is not interrupted during the switchover. Interface and kernel information is preserved. Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers. For more information about graceful restart, see “Graceful Restart Overview” on page 95.

Conditions That Trigger a Routing Engine Failover

The following events can result in an automatic change in Routing Engine mastership, depending on your configuration:

- The routing platform experiences a hardware failure. A change in Routing Engine mastership occurs if either the Routing Engine or the associated host module or subsystem is abruptly powered off. You can also configure the backup Routing Engine to take mastership if it detects a hard disk error on the master Routing Engine. To enable this feature, include the **failover on-disk-failure** statement at the [edit chassis redundancy] hierarchy level.

- The routing platform experiences a software failure, such as a kernel crash or a CPU lock. You must configure the backup Routing Engine to take mastership when it detects a loss of keepalive signal. To enable this failover method, include the `failover on-loss-of-keepalives` statement at the `[edit chassis redundancy]` hierarchy level.
- A specific software process fails. You can configure the backup Routing Engine to take mastership when one or more specified processes fail at least four times within 30 seconds. Include the `failover other-routing-engine` statement at the `[edit system processes process-name]` hierarchy level.

If any of these conditions is met, a message is logged and the backup Routing Engine attempts to take mastership. By default, an alarm is generated when the backup Routing Engine becomes active. After the backup Routing Engine takes mastership, it continues to function as master even after the originally configured master Routing Engine has successfully resumed operation. You must manually restore it to its previous backup status. (However, if at any time one of the Routing Engines is not present, the other Routing Engine becomes master automatically, regardless of how redundancy is configured.)

Default Routing Engine Redundancy Behavior

By default, JUNOS Software uses **re0** as the master Routing Engine and **re1** as the backup Routing Engine. Unless otherwise specified in the configuration, **re0** always becomes master when the acting master Routing Engine is rebooted.



NOTE: A single Routing Engine in the chassis always becomes the master Routing Engine even if it was previously the backup Routing Engine.

Perform the following steps to see how the default Routing Engine redundancy setting works:

1. Ensure that **re0** is the master Routing Engine.
2. Manually switch the state of Routing Engine mastership by issuing the `request chassis routing-engine master switch` command from the master Routing Engine. **re0** is now the backup Routing Engine and **re1** is the master Routing Engine.



NOTE: On the next reboot of the master Routing Engine, JUNOS Software returns the router to the default state because you have not configured the Routing Engines to maintain this state after a reboot.

3. Reboot the master Routing Engine **re1**.

The Routing Engine boots up and reads the configuration. Because you have not specified in the configuration which Routing Engine is the master, **re1** uses the default configuration as the backup. Now both **re0** and **re1** are in a backup state. JUNOS Software detects this conflict and, to prevent a no-master state, reverts to the default configuration to direct **re0** to become master.

Routing Engine Redundancy on a TX Matrix Router

On a routing matrix, all master Routing Engines in the TX Matrix router and connected T640 routers must run the same JUNOS Software release. Likewise, all backup Routing Engines in a routing matrix must run the same JUNOS Software release. When you run the same JUNOS Software release on all master and backup Routing Engines in the routing matrix, a change in mastership to any backup Routing Engine in the routing matrix does not cause a change in mastership in any other chassis in the routing matrix.

If the same JUNOS Software release is not running on all master and backup Routing Engines in the routing matrix, the following consequences occur when the **failover on-loss-of-keepalives** statement *is* included at the `[edit chassis redundancy]` hierarchy level:

- When the **failover on-loss-of-keepalives** statement is included at the `[edit chassis redundancy]` hierarchy level and you or a host subsystem initiates a change in mastership to the backup Routing Engine in the TX Matrix router, the master Routing Engines in the T640 routers detect a software release mismatch with the new master Routing Engine in the TX Matrix router and switch mastership to their backup Routing Engines.
- When you manually change mastership to a backup Routing Engine in a T640 router using the **request chassis routing-engine master** command, the new master Routing Engine in the T640 router detects a software release mismatch with the master Routing Engine in the TX Matrix router and relinquishes mastership to the original master Routing Engine. (Routing Engine mastership in the TX Matrix router does not switch in this case.)
- When a host subsystem initiates a change in mastership to a backup Routing Engine in a T640 router because the master Routing Engine has failed, the T640 router is logically disconnected from the TX Matrix router. To reconnect the T640 router, initiate a change in mastership to the backup Routing Engine in the TX Matrix router, or replace the failed Routing Engine in the T640 router and switch mastership to it. The replacement Routing Engine must be running the same software release as the master Routing Engine in the TX Matrix router.

If the same JUNOS Software release is not running on all master and backup Routing Engines in the routing matrix, the following consequences occur when the **failover on-loss-of-keepalives** statement *is not* included at the `[edit chassis redundancy]` hierarchy level:

- If you initiate a change in mastership to the backup Routing Engine in the TX Matrix router, all T640 routers are logically disconnected from the TX Matrix router. To reconnect the T640 routers, switch mastership of all master Routing Engines in the T640 routers to their backup Routing Engines.
- If you initiate a change in mastership to a backup Routing Engine in a T640 router, the T640 router is logically disconnected from the TX Matrix router. To reconnect the T640 router, switch mastership of the new master Routing Engine in the T640 router back to the original master Routing Engine.

Situations That Require You to Halt Routing Engines

Before you shut the power off to a routing platform that has two Routing Engines or before you remove the master Routing Engine, you must first halt the backup Routing Engine and then halt the master Routing Engine. Otherwise, you might need to reinstall JUNOS Software. You can use the **request system halt both-routing-engines** command on the master Routing Engine, which first shuts down the master Routing Engine and then shuts down the backup Routing Engine. To shut down only the backup Routing Engine, issue the **request system halt** command on the backup Routing Engine.

If you halt the master Routing Engine and do not power it off or remove it, the backup Routing Engine remains inactive unless you have configured it to become the master when it detects a loss of keepalive signal from the master Routing Engine.



NOTE: To restart the router, you must log in to the console port (rather than the Ethernet management port) of the Routing Engine. When you log in to the console port of the master Routing Engine, the system automatically reboots. After you log in to the console port of the backup Routing Engine, press Enter to reboot it.



NOTE: If you have upgraded the backup Routing Engine, first reboot it and then reboot the master Routing Engine.

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 chapter, 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 16
- Redundant FEBs on the M120 Router on page 16
- Redundant SSBs on the M20 Router on page 18
- Redundant SFMs on the M40e and M160 Routers on page 19

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 System Basics and Services Command Reference*.

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 System Basics Configuration Guide*. 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. 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 nonprimary-FEBs-to-backup), or a combination of both (one primary-to-backup and one or more nonprimary-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 a nonprimary FEB and a primary FEB is specified for the group, the backup FEB reboots so that the forwarding state from the nonprimary 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 following a switchover and this update can take minutes. If you do not want the interfaces to remain online during the switchover for a nonprimary 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 *JUNOS System Basics and Services Command Reference*.

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 *JUNOS System Basics and Services Command Reference*.

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 System Basics Configuration Guide*. 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 *JUNOS System Basics and Services Command Reference*.

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 *JUNOS System Basics and Services Command Reference*.

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 *JUNOS System Basics and Services Command Reference*.

Chapter 3

Routing Engine and Switching Control Board Redundancy Configuration Guidelines

This chapter includes the following topics:

- Chassis Redundancy Hierarchy on page 21
- Initial Routing Engine Configuration Example on page 22
- Copying a Configuration File from One Routing Engine to the Other on page 23
- Loading a Software Package from the Other Routing Engine on page 24
- Configuring Routing Engine Redundancy on page 24
- Configuring CFEB Redundancy on the M10i Router on page 28
- Configuring FEB Redundancy on the M120 Router on page 29
- Example: Configuring FEB Redundancy on page 30
- Configuring SFM Redundancy on M40e and M160 Routers on page 31
- Configuring SSB Redundancy on the M20 Router on page 31

Chassis Redundancy Hierarchy

The following redundancy statements are available at the [edit chassis] hierarchy level:

```
redundancy {  
  cfep slot (always | preferred);  
  failover on-disk-failure;  
  failover on-loss-of-keepalives;  
  feb {  
    redundancy-group group-name {  
      feb slot-number (backup | primary);  
      description description;  
      no-auto-failover;  
    }  
  }  
  graceful-switchover;  
  keepalive-time seconds;  
  routing-engine slot-number (master | backup | disabled);  
  sfm slot-number (always | preferred);
```

```

    ssb slot-number (always | preferred);
}

```

Initial Routing Engine Configuration Example

You can use configuration groups to ensure that the correct IP addresses are used for each Routing Engine and to maintain a single configuration file for both Routing Engines.

The following example defines configuration groups **re0** and **re1** with separate IP addresses. These well-known configuration group names take effect only on the appropriate Routing Engine.

```

groups {
  re0 {
    system {
      host-name my-re0;
    }
    interfaces {
      fxp0 {
        description "10/100 Management interface";
        unit 0 {
          family inet {
            address 10.255.2.40/24;
          }
        }
      }
    }
  }
  re1 {
    system {
      host-name my-re1;
    }
    interfaces {
      fxp0 {
        description "10/100 Management interface";
        unit 0 {
          family inet {
            address 10.255.2.41/24;
          }
        }
      }
    }
  }
}

```

You can assign an additional IP address to the **fxp0** (management) interface on both Routing Engines. The assigned address uses the **master-only** keyword and is identical for both Routing Engines, ensuring that the IP address for the master Routing Engine can be accessed at any time. The address is active only on the master Routing Engine's **fxp0** (management) interface. During a Routing Engine switchover, the address moves over to the new master Routing Engine.

For example, on **re0**, the configuration would be:

```
[edit groups re0 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.132/25;
  }
}
```

On **re1**, the configuration would be:

```
[edit groups re1 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.133/25;
  }
}
```

For more information about the initial configuration of dual Routing Engines, see the *JUNOS Software Installation and Upgrade Guide*. For more information about assigning an additional IP address to the **fxp0** (management) interface with the **master-only** keyword on both Routing Engines, see the *JUNOS CLI User Guide*.

Copying a Configuration File from One Routing Engine to the Other

You can use either the console port or the management Ethernet port to establish connectivity between the two Routing Engines. You can then copy or use FTP to transfer the configuration from the master to the backup, and load the file and commit it in the normal way.

To connect to the other Routing Engine using the management Ethernet port, issue the following command:

```
user@host> request routing-engine login (other-routing-engine | re0 | re1)
```

On a TX Matrix router, to make connections to the other Routing Engine using the management Ethernet port, issue the following command:

```
user@host> request routing-engine login (backup | lcc number | master | other-routing-engine | re0 | re1)
```

For more information about the **request routing-engine login** command, see the *JUNOS System Basics and Services Command Reference*.

To copy a configuration file from one Routing Engine to the other, issue the **file copy** command:

```
user@host> file copy source destination
```

In this case, *source* is the name of the configuration file. These files are stored in the directory `/config`. The active configuration is `/config/juniper.conf`, and older configurations are in `/config/juniper.conf {1...9}`. The *destination* is a file on the other Routing Engine.

The following example copies a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

The following example copies a configuration file from Routing Engine 0 to Routing Engine 1 on a TX Matrix router:

```
user@host> file copy /config/juniper.conf scc-re1:/var/tmp/copied-juniper.conf
```

To load the configuration file, enter the `load replace` command at the `[edit]` hierarchy level:

```
user@host> load replace /var/tmp/copied-juniper.conf
```



CAUTION: Make sure you change any IP addresses specified in `fxp0` on Routing Engine 0 to addresses appropriate for Routing Engine 1.

Loading a Software Package from the Other Routing Engine

You can load a package from the other Routing Engine onto the local Routing Engine using the existing `request system software add` *package-name* command:

```
user@host> request system software add re(0|1):/filename
```

In the *re* portion of the URL, specify the number of the other Routing Engine. In the *filename* portion of the URL, specify the path to the package. Packages are typically in the directory `/var/sw/pkg`.

Configuring Routing Engine Redundancy

The following sections describe how to configure Routing Engine redundancy:



NOTE: To complete the tasks in the following sections, `re0` and `re1` configuration groups must be defined. For more information about configuration groups, see the *JUNOS CLI User Guide*.

- Modifying the Default Routing Engine Mastership on page 25
- Configuring Automatic Failover to the Backup Routing Engine on page 25
- Manually Switching Routing Engine Mastership on page 27
- Verifying Routing Engine Redundancy Status on page 27

Modifying the Default Routing Engine Mastership

For routers with two Routing Engines, you can configure which Routing Engine is the master and which is the backup. By default, the Routing Engine in slot 0 is the master (**re0**) and the one in slot 1 is the backup (**re1**).



NOTE: In systems with two Routing Engines, both Routing Engines cannot be configured to be master at the same time. This configuration causes the commit check to fail.

To modify the default configuration, include the **routing-engine** statement at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
routing-engine slot-number (master | backup | disabled);
```

slot-number can be 0 or 1. To configure the Routing Engine to be the master, specify the **master** option. To configure it to be the backup, specify the **backup** option. To disable a Routing Engine, specify the **disabled** option.



NOTE: To switch between the master and the backup Routing Engines, you must modify the configuration and then activate it by issuing the **commit synchronize** command.

Configuring Automatic Failover to the Backup Routing Engine

The following sections describe how to configure automatic failover to the backup Routing Engine when certain failures occur on the master Routing Engine.

- Without Interruption to Packet Forwarding on page 25
- On Detection of a Hard Disk Error on the Master Routing Engine on page 26
- On Detection of a Loss of Keepalive Signal from the Master Routing Engine on page 26
- When a Software Process Fails on page 27

Without Interruption to Packet Forwarding

For routers with two Routing Engines, you can configure graceful Routing Engine switchover (GRES). When graceful switchover is configured, socket reconnection occurs seamlessly without interruption to packet forwarding. For information about how to configure graceful Routing Engine switchover, see “Graceful Routing Engine Switchover Configuration Guidelines” on page 53.

On Detection of a Hard Disk Error on the Master Routing Engine

After you configure a backup Routing Engine, you can direct it to take mastership automatically if it detects a hard disk error from the master Routing Engine. To enable this feature, include the `failover on-disk-failure` statement at the `[edit chassis redundancy]` hierarchy level.

```
[edit chassis redundancy]
failover on-disk-failure;
```

On Detection of a Loss of Keepalive Signal from the Master Routing Engine

After you configure a backup Routing Engine, you can direct it to take mastership automatically if it detects a loss of keepalive signal from the master Routing Engine.

To enable failover on receiving a loss of keepalive signal, include the `failover on-loss-of-keepalives` statement at the `[edit chassis redundancy]` hierarchy level:

```
[edit chassis redundancy]
failover on-loss-of-keepalives;
```

When graceful Routing Engine switchover is not configured, by default, failover occurs after 300 seconds (5 minutes). You can configure a shorter or longer time interval. To change the keepalive time period, include the `keepalive-time` statement at the `[edit chassis redundancy]` hierarchy level:

```
[edit chassis redundancy]
keepalive-time seconds;
```

The range for `keepalive-time` is 2 through 10,000 seconds.

The following example describes the sequence of events if you configure the backup Routing Engine to detect a loss of keepalive signal in the master Routing Engine:

1. Manually configure a `keepalive-time` of 25 seconds.
2. After the Packet Forwarding Engine connection to the primary Routing Engine is lost and the keepalive timer expires, packet forwarding is interrupted.
3. After 25 seconds of keepalive loss, a message is logged, and the backup Routing Engine attempts to take mastership. An alarm is generated when the backup Routing Engine becomes active, and the display is updated with the current status of the Routing Engine.
4. After the backup Routing Engine takes mastership, it continues to function as master.



NOTE: When graceful Routing Engine switchover is 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.

A former master Routing Engine becomes a backup Routing Engine if it returns to service after a failover to the backup Routing Engine. To restore master status to the former master Routing Engine, you can use the `request chassis routing-engine master switch` operational mode command.

If at any time one of the Routing Engines is not present, the remaining Routing Engine becomes master automatically, regardless of how redundancy is configured.

When a Software Process Fails

To configure automatic switchover to the backup Routing Engine if a software process fails, include the `failover other-routing-engine` statement at the `[edit system processes process-name]` hierarchy level:

```
[edit system processes] process-name
failover other-routing-engine;
```

process-name is one of the valid process names. If this statement is configured for a process, and that process fails four times within 30 seconds, the router reboots from the other Routing Engine. Another statement available at the `[edit system processes]` hierarchy level is `failover alternate-media`. For information about the alternate media option, see the *JUNOS System Basics Configuration Guide*.

Manually Switching Routing Engine Mastership

To manually switch the Routing Engine mastership, use one of the following commands:

- On the backup Routing Engine, request that the backup Routing Engine take mastership by issuing the `request chassis routing-engine master acquire` command.
- On the master Routing Engine, request that the backup Routing Engine take mastership by using the `request chassis routing-engine master release` command.
- On either Routing Engine, switch mastership by issuing the `request chassis routing-engine master switch` command.

Verifying Routing Engine Redundancy Status

A separate log file is provided for redundancy logging at `/var/log/mastership`. To view the log, use the `file show /var/log/mastership` command. Table 3 on page 27 lists the mastership log event codes and descriptions.

Table 3: Routing Engine Mastership Log

Event Code	Description
E_NULL = 0	The event is a null event.
E_CFG_M	The Routing Engine is configured as master.
E_CFG_B	The Routing Engine is configured as backup.

Table 3: Routing Engine Mastership Log *(continued)*

Event Code	Description
E_CFG_D	The Routing Engine is configured as disabled.
E_MAXTRY	The maximum number of tries to acquire or release mastership was exceeded.
E_REQ_C	A claim mastership request was sent.
E_ACK_C	A claim mastership acknowledgement was received.
E_NAK_C	A claim mastership request was not acknowledged.
E_REQ_Y	Confirmation of mastership is requested.
E_ACK_Y	Mastership is acknowledged.
E_NAK_Y	Mastership is not acknowledged.
E_REQ_G	A release mastership request was sent by a Routing Engine.
E_ACK_G	The Routing Engine acknowledged release of mastership.
E_CMD_A	The command <code>request chassis routing-engine master acquire</code> was issued from the backup Routing Engine.
E_CMD_F	The command <code>request chassis routing-engine master acquire force</code> was issued from the backup Routing Engine.
E_CMD_R	The command <code>request chassis routing-engine master release</code> was issued from the master Routing Engine.
E_CMD_S	The command <code>request chassis routing-engine master switch</code> was issued from a Routing Engine.
E_NO_ORE	No other Routing Engine is detected.
E_TMOUT	A request timed out.
E_NO_IPC	Routing Engine connection was lost.
E_ORE_M	Other Routing Engine state was changed to master.
E_ORE_B	Other Routing Engine state was changed to backup.
E_ORE_D	Other Routing Engine state was changed to disabled.

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, 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]
  cfep 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 cfep master switch** command. To view CFEB status, issue the **show chassis cfep** command.

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:

```
redundancy-group group-name {
  feb slot-number (backup | primary);
  description description;
  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 *JUNOS System Basics and Services Command Reference*.

Example: Configuring FEB Redundancy

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 System Basics Configuration Guide*.

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;
  }
}

```

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 *JUNOS System Basics and Services Command Reference*.

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 *JUNOS System Basics and Services Command Reference*.

Chapter 4

Summary of Routing Engine and Switching Control Board Redundancy Statements

This chapter provides a reference for each of the Routing Engine and switching control board redundancy configuration statements. The statements are organized alphabetically.

cfeb

Syntax	<code>cfeb slot-number</code> (always preferred);
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On M10i routers only, configure which Compact Forwarding Engine Board (CFEB) is the master and which is the backup.
Default	By default, the CFEB in slot 0 is the master and the CFEB in slot 1 is the backup.
Options	<code>slot-number</code> —Specify which slot is the master and which is the backup. <code>always</code> —Define this CFEB as the sole device. <code>preferred</code> —Define this CFEB as the preferred device of at least two.
Usage Guidelines	See “Configuring CFEB Redundancy on the M10i Router” on page 28.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.

description

Syntax	<code>description <i>description</i>;</code>
Hierarchy Level	<code>[edit chassis redundancy feb redundancy redundancy-group <i>name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Provide a description of the FEB redundancy group.
Options	<i>description</i> —Provide a description for the FEB redundancy group.
Usage Guidelines	See “Configuring FEB Redundancy on the M120 Router” on page 29
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.

failover on-disk-failure

Syntax	<code>failover on-disk-failure;</code>
Hierarchy Level	<code>[edit chassis redundancy]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Instruct the backup router to take mastership if it detects hard disk errors on the master Routing Engine.
Usage Guidelines	See “On Detection of a Hard Disk Error on the Master Routing Engine” on page 26.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.

failover on-loss-of-keepalives

Syntax	failover on-loss-of-keepalives;
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Instruct the backup router to take mastership if it detects a loss of keepalive signal from the master Routing Engine.
Default	<p>If the <code>failover on-loss-of-keepalives</code> statement at the [edit chassis redundancy] hierarchy level <i>is not</i> included and graceful Routing Engine switchover (GRES) is <i>not</i> enabled, failover cannot occur.</p> <p>When the <code>failover on-loss-of-keepalives</code> statement <i>is</i> included and graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the <code>failover on-loss-of-keepalives</code> statement <i>is</i> 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>
Usage Guidelines	See “On Detection of a Loss of Keepalive Signal from the Master Routing Engine” on page 26.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Topics	keepalive-time

failover other-routing-engine

Syntax	failover other-routing-engine;
Hierarchy Level	[edit system processes <i>process-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Instruct the backup Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, the router reboots from the backup Routing Engine.
Options	<i>process-name</i> —One of the valid software process names. A few examples are disk-monitoring, ethernet-link-fault-management, kernel-replication, redundancy-interface-process, and vrrp.
Usage Guidelines	See “When a Software Process Fails” on page 27.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Topics	For information about the failover alternate-media configuration statement at the [edit system processes <i>process-name</i>] hierarchy level, see the <i>JUNOS System Basics Configuration Guide</i> .

feb

- feb (Creating a Redundancy Group) on page 37
- feb (Assigning a FEB to a Redundancy Group) on page 37

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 Release 8.2.

Description On M120 routers only, configure a Forwarding Engine Board (FEB) redundancy group.

Options The remaining statements are described separately.

Usage Guidelines See “Configuring FEB Redundancy on the M120 Router” on page 29.

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

feb (Assigning a FEB to a Redundancy Group)

Syntax feb *slot-number* (backup | primary);

Hierarchy Level [edit chassis redundancy feb redundancy-group *group-name*]

Release Information Statement introduced in JUNOS Release 8.2.

Description On M120 routers only, configure a Forwarding Engine Board (FEB) as part of a FEB redundancy group.

Options *slot-number*—Slot number of the FEB. The range of values is from 0 to 5.

backup—(Optional) For each redundancy group, you must configure exactly one backup FEB.

primary—(Optional) For each redundancy group, you can optionally configure one primary FEB.

Usage Guidelines See “Configuring FEB Redundancy on the M120 Router” on page 29.

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

keepalive-time

Syntax	<code>keepalive-time seconds;</code>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the time period that must elapse before the backup router takes mastership when it detects loss of the keepalive signal.
Default	<p>If the <code>failover on-loss-of-keepalives</code> statement at the [edit chassis redundancy] hierarchy level <i>is not</i> included and graceful Routing Engine switchover (GRES) is <i>not</i> enabled, failover cannot occur.</p> <p>When the <code>failover on-loss-of-keepalives</code> statement <i>is</i> included and graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the <code>failover on-loss-of-keepalives</code> statement <i>is</i> 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>
Options	<code>seconds</code> —Time before the backup router takes mastership when it detects loss of the keepalive signal. The range of values is 2 through 10,000.
Usage Guidelines	See “On Detection of a Loss of Keepalive Signal from the Master Routing Engine” on page 26.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	failover on-loss-of-keepalives

no-auto-failover

Syntax	no-auto-failover;
Hierarchy Level	[edit chassis redundancy feb redundancy-group <i>group-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable automatic failover to a backup FEB when an active FEB in a redundancy group fails.
Default	Automatic failover is enabled by default.
Usage Guidelines	See “Configuring FEB Redundancy on the M120 Router” on page 29.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

redundancy

Syntax	<pre> redundancy { cfeb slot-number (always preferred); feb { redundancy-group group-name { description description; feb slot-number (backup primary); no-auto-failover; } } failover on-disk-failure; failover on-loss-of-keepalives; keepalive-time seconds; routing-engine slot-number (backup disabled master); sfm slot-number (always preferred); ssb slot-number (always preferred); } </pre>
Hierarchy Level	[edit chassis]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure redundancy options.
Options	The statements are explained separately.
Usage Guidelines	See “Routing Engine and Switching Control Board Redundancy Configuration Guidelines” on page 21.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

redundancy-group

Syntax	<code>redundancy-group <i>group-name</i> { description <i>description</i>; feb <i>slot-number</i> (backup primary); no-auto-failover; }</code>
Hierarchy Level	[edit chassis redundancy feb (Creating a Redundancy Group)]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	On M120 routers only, configure a Forwarding Engine Board (FEB) redundancy group.
Options	<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>
Usage Guidelines	See “Configuring FEB Redundancy on the M120 Router” on page 29.
Required Privilege Level	<p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>

routing-engine

Syntax	<code>routing-engine <i>slot-number</i> (backup disabled master);</code>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure Routing Engine redundancy.
Default	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.
Options	<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"> ■ master—Routing Engine in the specified slot is the master. ■ backup—Routing Engine in the specified slot is the backup. ■ disabled—Routing Engine in the specified slot is disabled.
Usage Guidelines	See “Configuring Routing Engine Redundancy” on page 24.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

sfm

Syntax	<code>sfm slot-number</code> (always preferred);
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On M40e and M160 routers, configure which Switching and Forwarding Module (SFM) is the master and which is the backup.
Default	By default, the SFM in slot 0 is the master and the SFM in slot 1 is the backup.
Options	<p><i>slot-number</i>—Specify which slot is the master and which is the backup. On the M40e router, <i>slot-number</i> can be 0 or 1. On the M160 router, <i>slot-number</i> can be 0 through 3.</p> <p><i>always</i>—Define this SFM as the sole device.</p> <p><i>preferred</i>—Define this SFM as the preferred device of at least two.</p>
Usage Guidelines	See “Configuring SFM Redundancy on M40e and M160 Routers” on page 31.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

ssb

Syntax	<code>ssb slot-number</code> (always preferred);
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On M20 routers, configure which System and Switch Board (SSB) is the master and which is the backup.
Default	By default, the SSB in slot 0 is the master and the SSB in slot 1 is the backup.
Options	<p><i>slot-number</i>—Specify which slot is the master and which is the backup.</p> <p><i>always</i>—Define this SSB as the sole device.</p> <p><i>preferred</i>—Define this SSB as the preferred device of at least two.</p>
Usage Guidelines	See “Configuring SSB Redundancy on the M20 Router” on page 31.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

Part 3

Graceful Routing Engine Switchover

- Graceful Routing Engine Switchover Overview on page 45
- Graceful Routing Engine Switchover Configuration Guidelines on page 53
- Summary of Graceful Routing Engine Switchover Configuration Statements on page 57

Chapter 5

Graceful Routing Engine Switchover Overview

This chapter contains the following sections:

- Understanding Graceful Routing Engine Switchover in the JUNOS Software on page 45
- Graceful Routing Engine Switchover System Requirements on page 48

Understanding Graceful Routing Engine Switchover in the JUNOS Software

This topic contains the following sections:

- Graceful Routing Engine Switchover Concepts on page 45
- Effects of a Routing Engine Switchover on page 48

Graceful Routing Engine Switchover Concepts

Graceful Routing Engine switchover (GRES) feature in JUNOS Software enables a routing platform with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. Graceful Routing Engine switchover preserves interface and kernel information. Traffic is not interrupted. However, graceful Routing Engine switchover does not preserve the control plane. Neighboring routers detect that the router has experienced a restart and react to the event in a manner prescribed by individual routing protocol specifications. To preserve routing during a switchover, graceful Routing Engine switchover must be combined with either graceful restart protocol extensions or nonstop active routing. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.



NOTE: To quickly restore or to preserve routing protocol state information during a switchover, graceful Routing Engine switchover must be combined with either graceful restart or nonstop active routing (NSR), respectively. For more information about graceful restart, see “Graceful Restart Overview” on page 95. For more information about nonstop active routing, see “Nonstop Active Routing Overview” on page 71.

If the backup Routing Engine does not receive a keepalive from the master Routing Engine after 2 seconds (4 seconds on M20 routers), it determines that the master Routing Engine has failed and takes mastership. The Packet Forwarding Engine seamlessly disconnects from the old master Routing Engine and reconnects to the new master Routing Engine. The Packet Forwarding Engine does not reboot, and traffic is not interrupted. The new master Routing Engine and the Packet Forwarding Engine then become synchronized. If the new master Routing Engine detects that the Packet Forwarding Engine state is not up to date, it resends state update messages.



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

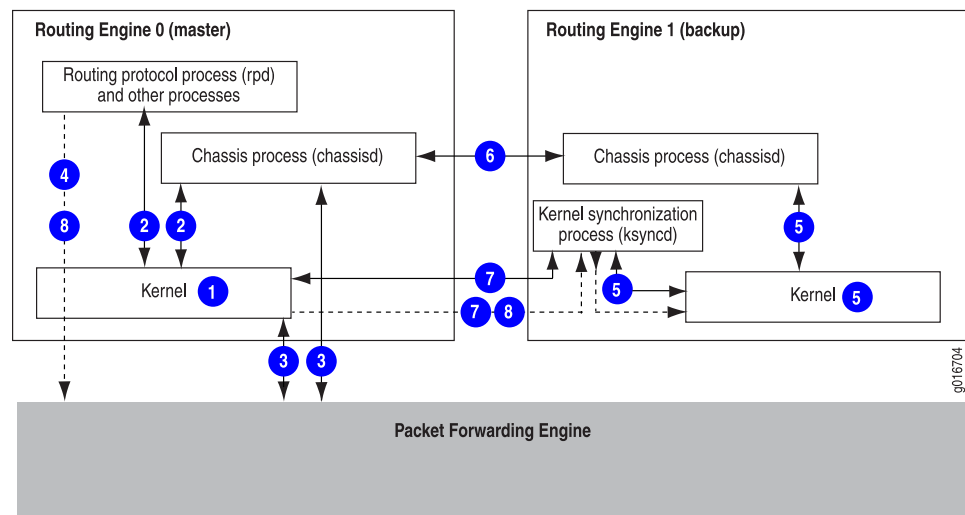
If the router displays a warning message similar to “Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset,” do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines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



NOTE: We do not recommend performing a commit operation on the backup Routing Engine when Graceful Routing Engine Switchover is enabled on the router.

Figure 1 on page 46 shows the system architecture of graceful Routing Engine switchover and the process a routing platform follows to prepare for a switchover.

Figure 1: Preparing for a Graceful Routing Engine Switchover

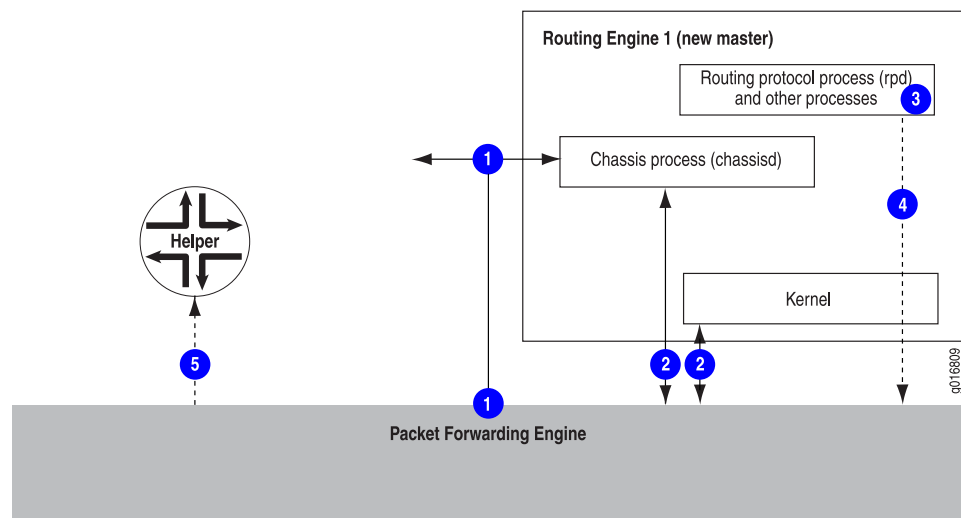


The switchover preparation process for graceful Routing Engine switchover follows these steps:

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

Figure 2 on page 47 shows the effects of a switchover on the routing platform.

Figure 2: Graceful Routing Engine Switchover Process



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

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

Effects of a Routing Engine Switchover

Table 4 on page 48 describes the effects of a Routing Engine switchover when no high availability features are enabled and when graceful Routing Engine switchover, graceful restart, and nonstop active routing features are enabled.

Table 4: Effects of a Routing Engine Switchover

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

Graceful Routing Engine Switchover System Requirements

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

- Graceful Routing Engine Switchover Platform Support on page 49
- Graceful Routing Engine Switchover Feature Support on page 49
- Graceful Routing Engine Switchover DPC Support on page 50
- Graceful Routing Engine Switchover and Subscriber Access on page 51
- Graceful Routing Engine Switchover PIC Support on page 51

Graceful Routing Engine Switchover Platform Support

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

- M20 and M40e routers—JUNOS Release 5.7 or later
- M10i router—JUNOS Release 6.1 or later
- M320 router—JUNOS Release 6.2 or later
- T320 router, T640 router, and TX Matrix router—JUNOS Release 7.0 or later
- M120 router—JUNOS Release 8.2 or later
- MX960 Ethernet Services router—JUNOS Release 8.3 or later
- MX480 Ethernet Services router—JUNOS Release 8.4 or later (8.4R2 recommended)
- MX240 Ethernet Services router—JUNOS Release 9.0 or later
- T1600 router—JUNOS Release 8.5 or later
- TX Matrix Plus router—JUNOS Release 9.6 or later

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

Graceful Routing Engine Switchover Feature Support

Graceful Routing Engine switchover supports most JUNOS Software features in Release 5.7 and later. Particular JUNOS Software features require specific versions of JUNOS Software. See Table 5 on page 49.

Table 5: Graceful Routing Engine Switchover Feature Support

Application	JUNOS Software Release
Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP) and aggregated SONET interfaces	6.2
Asynchronous Transfer Mode (ATM) virtual circuits (VCs)	6.2
Logical systems	6.3
NOTE: Beginning with JUNOS Software Release 9.3, the logical router feature has been renamed logical system.	
Multicast	6.4 (7.0 for TX Matrix router)
Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR)	7.0

Table 5: Graceful Routing Engine Switchover Feature Support (*continued*)

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

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

- When graceful Routing Engine switchover and aggregated Ethernet interfaces are configured in the same system, the aggregated Ethernet interfaces must not be configured for fast-polling LACP. When fast polling is configured, the LACP polls time out at the remote end during the Routing Engine mastership switchover. When LACP polling times out, the aggregated link and interface are disabled. The Routing Engine mastership change is fast enough that standard and slow LACP polling do not time out during the procedure.
- VRRP changes mastership when a Routing Engine switchover occurs, even when graceful Routing Engine switchover is configured.
- The extended DHCP *local* server does not support graceful Routing Engine switchover. Address-assignment pools do not support graceful Routing Engine switchover. For more information, see the *JUNOS System Basics Configuration Guide*.

Graceful Routing Engine Switchover DPC Support

Graceful Routing Engine switchover supports all Dense Port Concentrators (DPCs) on the MX Series Ethernet Services Routers running the appropriate version of JUNOS Software. For more information about DPCs, see the *MX Series DPC Guide*.

Graceful Routing Engine Switchover and Subscriber Access

Graceful Routing Engine switchover currently supports most of the features directly associated with dynamic DHCP subscriber access. However, Graceful Routing Engine switchover does not support unified in-service software upgrade (ISSU) on systems that are running DHCP subscriber access. In addition, any partially complete subscriber secure policy activations that are lost during a failure are not recovered.

Graceful Routing Engine Switchover PIC Support

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

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

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



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

Chapter 6

Graceful Routing Engine Switchover Configuration Guidelines

This chapter contains the following information:

- Configuring Graceful Routing Engine Switchover on page 53
- Requirements for Routers with a Backup Router Configuration on page 54
- Resetting Local Statistics on page 55

Configuring Graceful Routing Engine Switchover

This section contains the following topics:

- Enabling Graceful Routing Engine Switchover on page 53
- Synchronizing the Routing Engine Configuration on page 54
- Verifying Graceful Routing Engine Switchover Operation on page 54

Enabling Graceful Routing Engine Switchover

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

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

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

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

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

Synchronizing the Routing Engine Configuration



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

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

Verifying Graceful Routing Engine Switchover Operation

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

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



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

For more information about the `show system switchover` command, see the *JUNOS System Basics and Services Command Reference*.

Requirements for Routers with a Backup Router Configuration

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



NOTE: If you have a configuration in which multiple static routes point to a gateway from `fxp0`, you must either configure specific prefixes for the static routes or include the `retain` flag at the `[edit routing-options static route]` hierarchy level.

For example, if you configure the static route `172.16.0.0/12` from `fxp0` for management purposes, you must specify the backup router configuration as follows:

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

Resetting Local Statistics

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

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. After a graceful Routing Engine switchover, we recommend that you issue the `clear interface statistics` (`interface-name` | `all`) command to reset the cumulative values for local statistics. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the `clear` command to clear statistics and protocol database information, see the *JUNOS System Basics and Services Command Reference*.



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

Chapter 7

Summary of Graceful Routing Engine Switchover Configuration Statements

This chapter provides a reference for the graceful Routing Engine switchover configuration statement.

graceful-switchover

Syntax	graceful-switchover;
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.
Usage Guidelines	See “Configuring Graceful Routing Engine Switchover” on page 53.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

Part 4

Nonstop Bridging

- Nonstop Bridging Overview on page 61
- Nonstop Bridging Configuration Guidelines on page 65
- Summary of Nonstop Bridging Statements on page 67

Chapter 8

Nonstop Bridging Overview

This chapter contains the following information:

- Nonstop Bridging Concepts on page 61
- Nonstop Bridging System Requirements on page 63

Nonstop Bridging Concepts

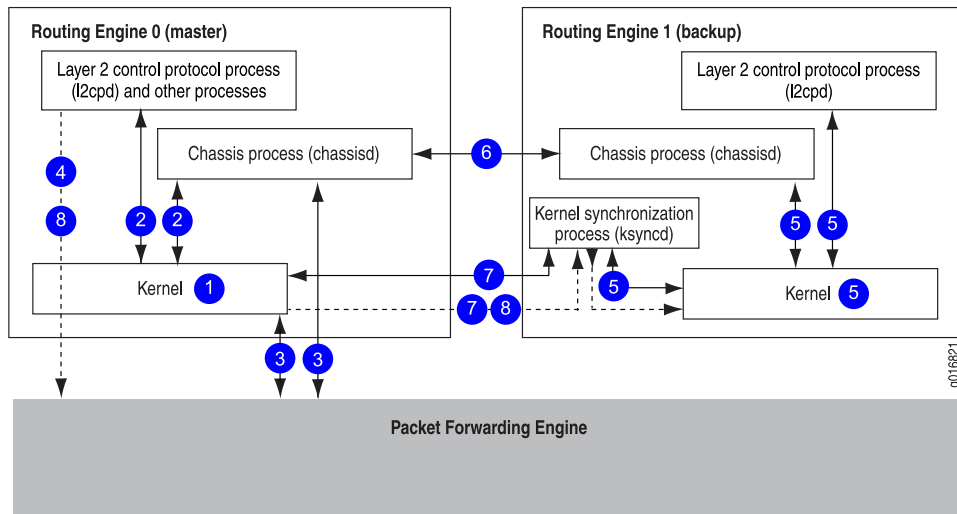
Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop bridging also saves Layer 2 Control Protocol (L2CP) information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine.



NOTE: To use nonstop bridging, you must first enable graceful Routing Engine switchover on your routing platform. For more information about graceful Routing Engine switchover, see “Graceful Routing Engine Switchover Overview” on page 45.

Figure 3 on page 62 shows the system architecture of nonstop bridging and the process a routing platform follows to prepare for a switchover.

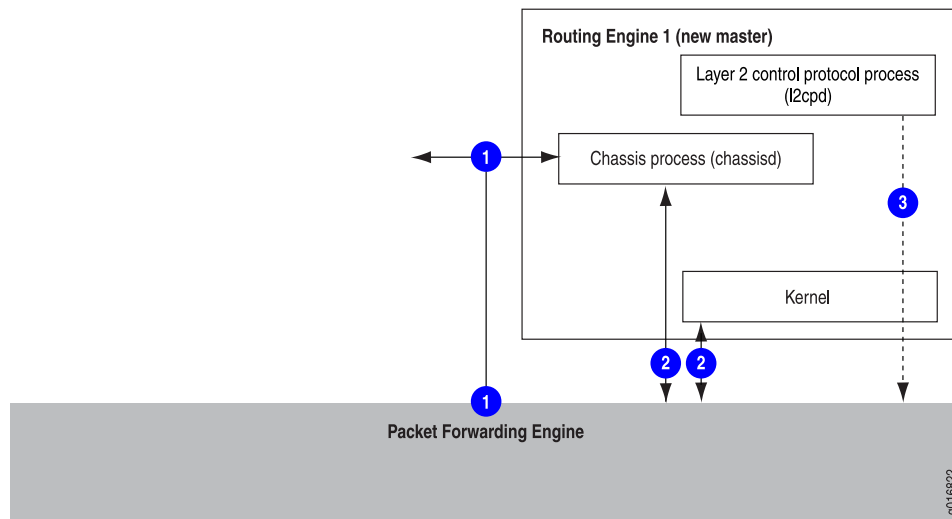
Figure 3: Nonstop Bridging Switchover Preparation Process



The switchover preparation process for nonstop bridging follows these steps:

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

Figure 4 on page 63 shows the effects of a switchover on the routing platform.

Figure 4: Nonstop Bridging During a Switchover

The switchover process follows these steps:

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

Nonstop Bridging System Requirements

This topic contains the following sections:

- Platform Support on page 63
- Protocol Support on page 64

Platform Support

Nonstop bridging is supported on MX Series Ethernet Services Routers. Your system must be running JUNOS Release 8.4 or later.



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

Protocol Support

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

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

Chapter 9

Nonstop Bridging Configuration Guidelines

This chapter includes the following topics:

- Configuring Nonstop Bridging on page 65

Configuring Nonstop Bridging

This section includes the following topics:

- Enabling Nonstop Bridging on page 65
- Synchronizing the Routing Engine Configuration on page 65
- Verifying Nonstop Bridging Operation on page 66

Enabling Nonstop Bridging

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

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

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

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

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

Synchronizing the Routing Engine Configuration

When you configure nonstop bridging, you must also include the `commit synchronize` statement at the `[edit system]` hierarchy level so that, by default, when you issue the `commit` command, the configuration changes are synchronized on both Routing Engines. If you issue the `commit synchronize` command at the `[edit]` hierarchy level

on the backup Routing Engine, the JUNOS system software displays a warning and commits the candidate configuration.



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

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

Verifying Nonstop Bridging Operation

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

Chapter 10

Summary of Nonstop Bridging Statements

This chapter provides a reference for the `nonstop-bridging` configuration statement.

nonstop-bridging

Syntax	<code>nonstop-bridging;</code>
Hierarchy Level	[edit protocols layer2-control]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.
Usage Guidelines	See “Configuring Nonstop Bridging” on page 65.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

Part 5

Nonstop Active Routing

- Nonstop Active Routing Overview on page 71
- Nonstop Active Routing Configuration Guidelines on page 81
- Summary of Nonstop Active Routing Configuration Statements on page 87

Chapter 11

Nonstop Active Routing Overview

This chapter contains the following information:

- Nonstop Active Routing Concepts on page 71
- Nonstop Active Routing System Requirements on page 73

Nonstop Active Routing Concepts

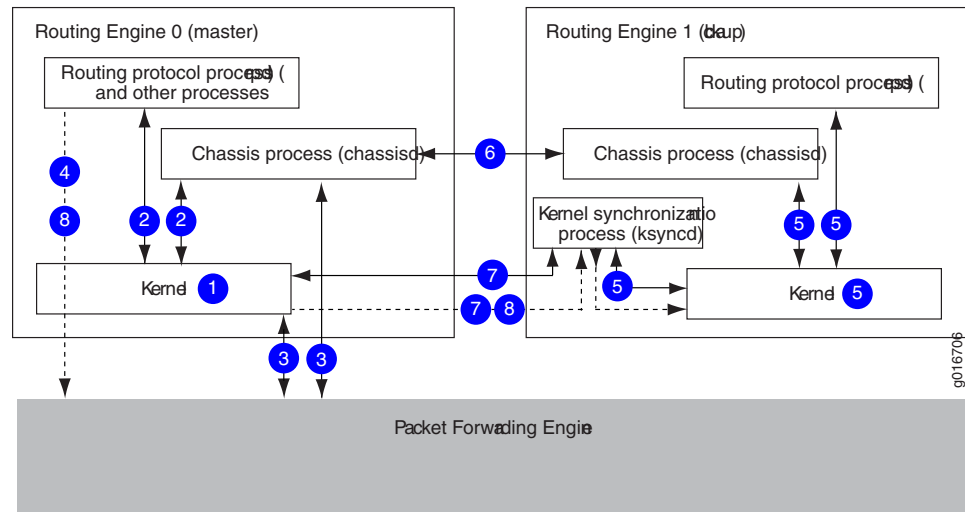
Nonstop active routing (NSR) uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop active routing also saves routing protocol information by running the routing protocol process (rpd) on the backup Routing Engine. By saving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



NOTE: To use nonstop active routing, you must first enable graceful Routing Engine switchover on your routing platform. For more information about graceful Routing Engine switchover, see “Graceful Routing Engine Switchover Overview” on page 45.

Figure 5 on page 72 shows the system architecture of nonstop active routing and the process a routing platform follows to prepare for a switchover.

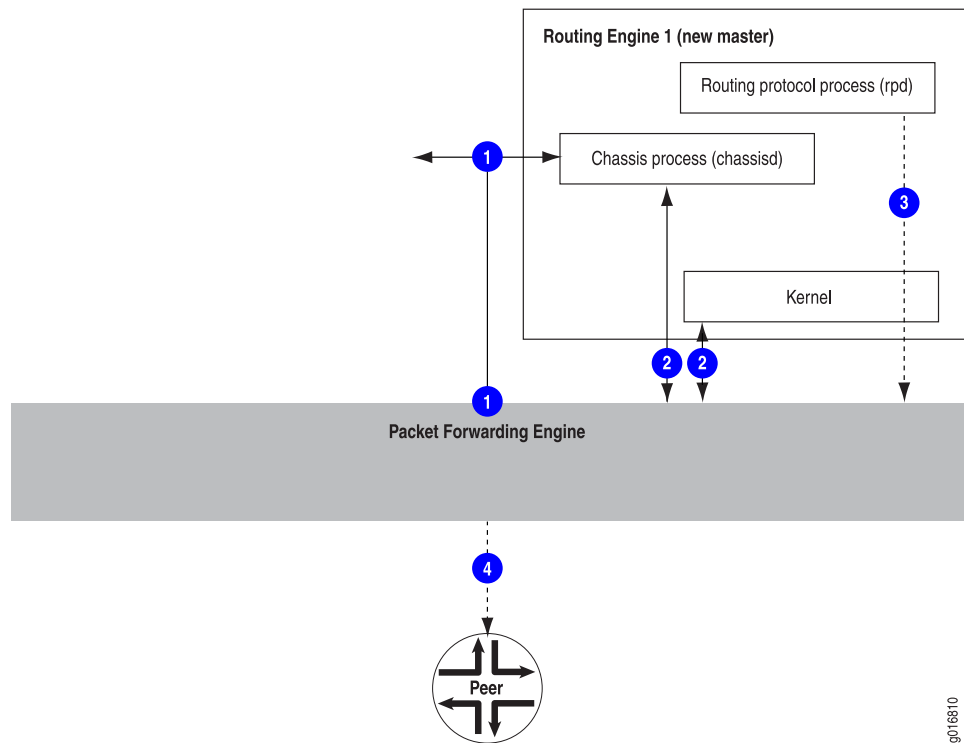
Figure 5: Nonstop Active Routing Switchover Preparation Process



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

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

Figure 6 on page 73 shows the effects of a switchover on the routing platform.

Figure 6: Nonstop Active Routing During a Switchover

The switchover process follows these steps:

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

Nonstop Active Routing System Requirements

This section contains the following topics:

- Nonstop Active Routing Platform Support on page 74
- Nonstop Active Routing Protocol and Feature Support on page 74
- Nonstop Active Routing BFD Support on page 76
- Nonstop Active Routing BGP Support on page 76
- Nonstop Active Routing Layer 2 Circuit and LDP-Based VPLS Support on page 77

- Nonstop Active Routing PIM Support on page 77
- Nonstop Active Routing Support for RSVP-TE LSPs on page 79

Nonstop Active Routing Platform Support

Table 6 on page 74 lists the platforms that support nonstop active routing.

Table 6: Nonstop Active Routing Platform Support

Platform	JUNOS Software Release
M10i router	8.4 or later
M20 router	8.4 or later
M40e router	8.4 or later
M120 router	9.0 or later
M320 router	8.4 or later
MX Series routers	9.0 or later
T320 router, T640 router, and TX Matrix router	8.4 or later
T1600 router	8.5 or later
TX Plus Matrix router	10.0 or later



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

Nonstop Active Routing Protocol and Feature Support

Table 7 on page 74 lists the protocols that are supported by nonstop active routing.

Table 7: Nonstop Active Routing Protocol and Feature Support

Protocol	JUNOS Software Release
Bidirectional forwarding detection (BFD)	8.5 or later
For more information about BFD support, see “Nonstop Active Routing BFD Support” on page 76.	
BGP	8.4 or later
For more information about nonstop active routing support for BGP, see “Nonstop Active Routing BGP Support” on page 76.	

Table 7: Nonstop Active Routing Protocol and Feature Support (*continued*)

Protocol	JUNOS Software Release
IS-IS	8.4 or later
LDP	8.4 or later
LDP-based virtual private LAN service (VPLS)	9.3 or later
LDP OAM (operation, administration, and management) features	9.6 or later
Layer 2 circuits	9.2 or later
Layer 2 VPNs	9.1 or later
Layer 3 VPNs (see the first Note after this table for restrictions)	9.2 or later
OSPF/OSPFv3	8.4 or later
Protocol Independent Multicast (PIM) (for IPv4)	9.3 or later
For more information about nonstop active routing support for PIM, see “Nonstop Active Routing PIM Support” on page 77.	
RIP and RIP next generation (RIPng)	9.0 or later
RSVP-TE LSP	9.5 or later
For more information about nonstop active routing support for RSVP-TE LSPs, see “Nonstop Active Routing Support for RSVP-TE LSPs” on page 79.	
VPLS	9.1 or later



NOTE: Layer 3 VPN support does not include RSVP-based tunnels, dynamic GRE tunnels, multicast VPNs, or BGP flow routes. The following OSPF features and configuration statements are also not supported:

- `domain-id` *domain-id* statement at the [edit routing-instances *routing-instance-name* protocols (ospf | ospf3)] hierarchy level
- `domain-vpn-tag` *number* statement at the [edit routing-instances *routing-instance-name* protocols (ospf | ospf3)] hierarchy level
- `metric` *number* statement at the [edit routing-instances *routing-instance-name* protocols ospf area *area-id* sham-link-remote] hierarchy level
- `sham-link local` *address* statement at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level
- `sham-link-remote` *address* <metric *number*> statement at the [edit routing-instances *routing-instance-name* protocols ospf area *area-id*] hierarchy level



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



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

Nonstop Active Routing BFD Support

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



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

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



NOTE: For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, the value for the `minimum-interval` configuration statement (a BFD liveness detection parameter) must be at least 2500 ms for Routing Engine-based sessions and at least 10 ms for distributed BFD sessions.

Nonstop Active Routing BGP Support

If the BGP peer in the master Routing Engine has negotiated address-family capabilities that are not supported for nonstop active routing, then the corresponding BGP neighbor state on the backup Routing Engine shows as idle. On switchover, the BGP session is reestablished from the new master Routing Engine.

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



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

- inet unicast
- inet labeled-unicast
- inet multicast
- inet6 labeled-unicast
- inet6 multicast
- inet6 unicast
- route-target
- l2vpn signaling
- inet6-vpn unicast
- inet-vpn unicast

Nonstop Active Routing Layer 2 Circuit and LDP-Based VPLS Support

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

Starting with Release 9.6, JUNOS Software extends nonstop active routing support to the Layer 2 circuit and LDP-based VPLS pseudowire redundant configurations.

Nonstop Active Routing PIM Support

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

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



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

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

Supported features:

- Auto-RP
- BFD
- Bootstrap router
- Dense mode
- Sparse mode (except for some subordinate features mentioned in the following list of unsupported features)
- Source-specific multicast (SSM)
- Static RPs

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

- Internet Group Management Protocol (IGMP) exclude mode
- IGMP snooping
- PIM for IPv6 and related features such as embedded RP and Multicast Listener Discovery (MLD)
- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies.
- Upstream assert synchronization

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

- Anycast RP
- Draft-Rosen multicast VPNs (MVPNs)
- Local RP
- Next-generation MVPNs with PIM provider tunnels
- PIM join load balancing

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

Nonstop Active Routing Support for RSVP-TE LSPs

JUNOS Software extends nonstop active routing support to transit label-switching routers (LSR) that are part of an RSVP-TE LSP. Nonstop active routing support on transit LSRs ensures that the master to backup Routing Engine switchover on an LSR remains transparent to the network neighbors and that the LSP information remains unaltered during and after the switchover. You can use the **show rsvp version** command to view the nonstop active routing mode and state on an LSR.

However, JUNOS Software does not support the following features for nonstop active routing on RSVP-TE transit and egress LSRs:

- Point-to-multipoint LSPs
- Generalized Multiprotocol Label Switching (GMPLS) and LSP hierarchy
- Interdomain or loose-hop expansion LSPs
- Ingress features such as circuit cross-connect (CCC), LSP advertising, and LDP tunneling are not supported. Similarly, ingress bypass LSPs are also not supported for nonstop active routing.

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

- Control plane statistics corresponding to the **show rsvp statistics** and **show rsvp interface detail | extensive** commands are not maintained across Routing Engine switchovers.
- Statistics from the backup Routing Engine are not reported for **show mpls lsp statistics** and **monitor mpls label-switched-path** commands. However, if a switchover occurs, the backup Routing Engine, after taking over as the master, starts reporting statistics. Note that the **clear statistics** command issued on the old master Routing Engine does not have any effect on the new master Routing Engine, which reports statistics, including any uncleared statistics.
- State timeouts may take additional time during nonstop active routing switchover. For example, if a switchover occurs after a neighbor has missed sending two hello messages to the master, the new master Routing Engine waits for another three hello periods before timing out the neighbor.
- When nonstop active routing is enabled, graceful restart is not supported. However, graceful restart helper mode is supported.

Chapter 12

Nonstop Active Routing Configuration Guidelines

This chapter contains the following sections:

- Configuring Nonstop Active Routing on page 81
- Tracing Nonstop Active Routing Synchronization Events on page 83
- Resetting Local Statistics on page 84
- Example: Configuring Nonstop Active Routing on page 84

Configuring Nonstop Active Routing

This section includes the following topics:

- Enabling Nonstop Active Routing on page 81
- Synchronizing the Routing Engine Configuration on page 82
- Verifying Nonstop Active Routing Operation on page 82

Enabling Nonstop Active Routing

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

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

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

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

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



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

To enable the routing platform to switch over to the backup Routing Engine when the routing protocol process (rpd) fails rapidly three times in succession, include the `other-routing-engine` statement at the `[edit system processes routing failover]` hierarchy level.

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

Synchronizing the Routing Engine Configuration

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

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

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

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

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



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

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

Verifying Nonstop Active Routing Operation

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

For more information on these commands, see the *JUNOS System Basics and Services Command Reference* and *JUNOS Routing Protocols and Policies Command Reference*, respectively.

When you enable nonstop active routing and issue routing-related operational mode commands on the backup Routing Engine (such as `show route`, `show bgp neighbor`, `show ospf database`, and so on), the output might not match the output of the same commands issued on the master Routing Engine.

To display BFD state replication status, issue the `show bfd session` command. The **replicated** flag appears in the output for this command when a BFD session has been replicated to the backup Routing Engine. For more information, see the *JUNOS Routing Protocols and Policies Command Reference*.

Tracing Nonstop Active Routing Synchronization Events

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

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

```
[edit protocols]
  bgp {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  isis {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  ldp {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  (ospf | ospf3) {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  (rip | ripng) {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  pim {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
```

```
    }
}
```

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

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

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

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

Resetting Local Statistics

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

Example: Configuring Nonstop Active Routing

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

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

```

    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.2.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.3.1.1/30;
      }
      family iso;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.2.1/32;
      }
      family iso {
        address 49.0004.1921.6800.2001.00;
      }
    }
  }
}
routing-options {
  nonstop-routing; # This enables nonstop active routing on the routing platform.
  router-id 192.168.2.1;
  autonomous-system 65432;
}
protocols {
  bgp {
    traceoptions {
      flag nsr-synchronization detail; # This logs nonstop active routing
      # events for BGP.
    }
    local-address 192.168.2.1;
    group external-group {
      type external;
      export BGP_export;
      neighbor 192.168.1.1 {
        family inet {
          unicast;
        }
      }
    }
  }
}

```

```

        peer-as 65103;
    }
}
group internal-group {
    type internal;
    neighbor 192.168.10.1;
    neighbor 192.168.11.1;
    neighbor 192.168.12.1;
}
}
isis {
    traceoptions {
        flag nsr-synchronization detail; # This logs nonstop active routing events
        # for IS-IS.
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        passive;
    }
}
ospf {
    traceoptions {
        flag nsr-synchronization detail; # This logs nonstop active routing events
        # for OSPF.
    }
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
}
}
policy-options {
    policy-statement BGP_export {
        term direct {
            from {
                protocol direct;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
}
}

```

Chapter 13

Summary of Nonstop Active Routing Configuration Statements

This chapter provides a reference for each of the nonstop active routing configuration statements. The statements are organized alphabetically.

commit synchronize

Syntax commit synchronize;
 <and-quit>
 <comment>
 <and-force>

Hierarchy Level [edit system]

Release Information Statement introduced in JUNOS Release 7.4.
 and-quit, comment, and and-force options added in JUNOS Release 8.5.

Description Configure the commit command to automatically result in a commit synchronize action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the commit command (requesting Routing Engine) copies and loads its candidate configuration to the other (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.

Synchronization only occurs between the Routing Engines within the same chassis. On the TX Matrix router, 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.



NOTE: When you configure nonstop active routing (NSR), you must include the commit synchronize statement. Otherwise, the commit fails.

Options and-quit—(Optional) Quit configuration mode if the commit synchronization succeeds.
 comment—(Optional) Write a message to the commit log.
 and-force—(Optional) Force a commit synchronization on the other Routing Engine (ignore warnings).

Usage Guidelines See “Synchronizing the Routing Engine Configuration” on page 82.

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

nonstop-routing

Syntax	nonstop-routing;
Hierarchy Level	[edit routing-options]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve routing protocol information.
Usage Guidelines	See “Configuring Nonstop Active Routing” on page 81.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

traceoptions

Syntax	<pre> traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> > <(world-readable no-world-readable)>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<pre> [edit protocols bfd], [edit protocols bgp], [edit protocols isis], [edit protocols ldp], [edit protocols ospf], [edit protocols ospf3], [edit protocols pim], [edit protocols rip], [edit protocols ripng], [edit routing-options] </pre>
Release Information	<p>Statement introduced before JUNOS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in JUNOS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in JUNOS Release 8.5.</p> <p>nsr-synchronization flag for RIP and RIPng added in JUNOS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in JUNOS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in JUNOS Release 9.3.</p>
Description	<p>Define tracing operations that track nonstop active routing (NSR) functionality in the router.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p>

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The nonstop active routing tracing options are:

- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing

flag-modifier—(Optional) Modifier for the tracing flag. Except for BFD sessions, you can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

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.

Usage Guidelines See “Tracing Nonstop Active Routing Synchronization Events” on page 83.

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

Part 6

Graceful Restart

- Graceful Restart Overview on page 95
- Graceful Restart Configuration Guidelines on page 103
- Summary of Graceful Restart Configuration Statements on page 141

Chapter 14

Graceful Restart Overview

This chapter contains the following sections:

- Graceful Restart Concepts on page 95
- Graceful Restart System Requirements on page 96
- Aggregate and Static Routes on page 96
- Graceful Restart and Routing Protocols on page 96
- Graceful Restart and MPLS-Related Protocols on page 98
- Graceful Restart and Layer 2 and Layer 3 VPNs on page 100
- Graceful Restart on Logical Systems on page 101

Graceful Restart Concepts

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. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

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

- Graceful restart for aggregate and static routes and for routing protocols—Provides protection for aggregate and static routes and for Border Gateway Protocol (BGP), End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), next-generation RIP (RIPng), and Protocol Independent Multicast (PIM) sparse mode routing protocols.
- Graceful restart for MPLS-related protocols—Provides protection for Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), circuit cross-connect (CCC), and translational cross-connect (TCC).
- Graceful restart for virtual private networks (VPNs)—Provides protection for Layer 2 and Layer 3 VPNs.

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

Most graceful restart implementations define two types of routers—the restarting router and the helper router. The restarting router requires rapid restoration of forwarding state information so it can resume the forwarding of network traffic. The helper router assists the restarting router in this process. Graceful restart configuration statements typically affect either the restarting router or the helper router.

Graceful Restart System Requirements

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

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

Aggregate and Static Routes

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

Graceful Restart and Routing Protocols

This section covers the following topics:

- BGP on page 97
- ES-IS on page 97

- IS-IS on page 97
- OSPF and OSPFv3 on page 98
- PIM Sparse Mode on page 98
- RIP and RIPng on page 98

BGP

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

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

ES-IS

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



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

IS-IS

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

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

OSPF and OSPFv3

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

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

PIM Sparse Mode

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

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

RIP and RIPng

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

Graceful Restart and MPLS-Related Protocols

This section contains the following topics:

- LDP on page 99
- RSVP on page 99
- CCC and TCC on page 99

LDP

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

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

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

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

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

RSVP

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

CCC and TCC

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

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

Graceful Restart and Layer 2 and Layer 3 VPNs

VPN graceful restart uses three types of restart functionality:

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

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

- The router must wait to receive all BGP NLRI information from other PE routers before advertising routes to the CE routers.
- The router must wait for all protocols in all routing instances to converge (or complete the restart process) before it sends CE router information to other PE routers. In other words, the router must wait for all instance information (whether derived from local configuration or advertisements received from a remote peer) to be processed before it sends this information to other PE routers.
- The router must preserve all forwarding state in the `instance.mpls.0` tables until the new labels and transit routes are allocated and announced to other PE routers (and CE routers in a carrier-of-carriers scenario).

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

Graceful Restart on Logical Systems

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

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

Chapter 15

Graceful Restart Configuration Guidelines

To implement graceful restart, you must perform the configuration tasks described in the following sections:

- Configuring Graceful Restart for Aggregate and Static Routes on page 103
- Configuring Routing Protocols Graceful Restart on page 103
- Configuring Graceful Restart for MPLS-Related Protocols on page 108
- Configuring VPN Graceful Restart on page 111
- Configuring Logical System Graceful Restart on page 112
- Verifying Graceful Restart Operation on page 113
- Example: Configuring Graceful Restart on page 115

Configuring Graceful Restart for Aggregate and Static Routes

To configure graceful restart for aggregate and static routes, include the `graceful-restart` statement at the `[edit routing-options]` hierarchy level. To disable graceful restart, include the `disable` statement at the `[edit routing-options graceful-restart]` hierarchy level.

Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- Configuring Graceful Restart Globally on page 104
- Configuring Graceful Restart Options for BGP on page 104
- Configuring Graceful Restart Options for ES-IS on page 105
- Configuring Graceful Restart Options for IS-IS on page 105
- Configuring Graceful Restart Options for OSPF and OSPFv3 on page 106
- Configuring Graceful Restart Options for RIP and RIPng on page 107
- Configuring Graceful Restart Options for PIM Sparse Mode on page 107
- Tracking Graceful Restart Events on page 108

Configuring Graceful Restart Globally

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



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

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

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

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



NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

Configuring Graceful Restart Options for BGP

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

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

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



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

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

Configuring Graceful Restart Options for ES-IS

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

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

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

Configuring Graceful Restart Options for IS-IS

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

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

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



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



NOTE: You can also track graceful restart events with the `traceoptions` statement at the `[edit protocols isis]` hierarchy level. For more information, see “Tracking Graceful Restart Events” on page 108.

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the `restart-duration` statement at the `[edit protocols (ospf | ospf3) graceful-restart]` hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the `notify-duration` at the `[edit protocols (ospf | ospf3) graceful-restart]` hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the `no-strict-lsa-checking` statement at the `[edit protocols (ospf | ospf3) graceful-restart]` hierarchy level.

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

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



NOTE: You can also track graceful restart events with the `traceoptions` statement at the `[edit protocols (ospf | ospf3)]` hierarchy level. For more information, see “Tracking Graceful Restart Events” on page 108.



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



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

Configuring Graceful Restart Options for RIP and RIPng

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

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

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

Configuring Graceful Restart Options for PIM Sparse Mode

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

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

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

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

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

```
[edit]
protocols {
  pim {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
```

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



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

Tracking Graceful Restart Events

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

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

Configuring Graceful Restart for MPLS-Related Protocols

This section contains the following topics:

- Configuring Graceful Restart Globally on page 109
- Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 109
- Configuring Graceful Restart Options for LDP on page 110

Configuring Graceful Restart Globally

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

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

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

Configuring Graceful Restart Options for RSVP, CCC, and TCC

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

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

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

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

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

Configuring Graceful Restart Options for LDP

You can configure the following optional statements at the [edit protocols ldp graceful-restart] hierarchy level:

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

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

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



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

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

Configuring VPN Graceful Restart

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

- Configuring Graceful Restart Globally on page 111
- Configuring Graceful Restart for the Routing Instance on page 111

Configuring Graceful Restart Globally

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

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

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

Configuring Graceful Restart for the Routing Instance

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

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

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

Configuring Logical System Graceful Restart

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

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

- Configuring Graceful Restart Globally on page 112
- Configuring Graceful Restart for a Routing Instance on page 112

Configuring Graceful Restart Globally

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

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

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

Configuring Graceful Restart for a Routing Instance

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

```
[edit]
logical-systems {
  logical-system-name {
    routing-instances {
      instance-name {
        routing-options {
```

```

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

```

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

Verifying Graceful Restart Operation

This topic contains the following sections:

- Graceful Restart Operational Mode Commands on page 113
- Verifying BGP Graceful Restart on page 113
- Verifying IS-IS and OSPF Graceful Restart on page 114
- Verifying CCC and TCC Graceful Restart on page 114

Graceful Restart Operational Mode Commands

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

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

For more information about these commands and a description of their output fields, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying BGP Graceful Restart

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

```
user@PE1> show bgp neighbor 192.255.10.1
```

```

Peer: 192.255.10.1+179 AS 64595 Local: 192.255.5.1+1106 AS 64595
Type: Internal State: Established Flags: <>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Export: [ static ]
Options:<Preference LocalAddress HoldTime GracefulRestart Damping PeerAS Refresh>

Local Address: 192.255.5.1 Holdtime: 90 Preference: 170
IPSec SA Name: hope
Number of flaps: 0
Peer ID: 192.255.10.1 Local ID: 192.255.5.1 Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 180
Stale routes from peer are kept for: 180
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast
NLRI that peer saved forwarding for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Table inet.0 Bit: 10000
RIB State: restart is complete
Send state: in sync
Active prefixes: 0
Received prefixes: 0
Suppressed due to damping: 0
Last traffic (seconds): Received 19 Sent 19 Checked 19
Input messages: Total 2 Updates 1 Refreshes 0 Octets 42
Output messages: Total 3 Updates 0 Refreshes 0 Octets 116
Output Queue[0]: 0

```

Verifying IS-IS and OSPF Graceful Restart

To view graceful restart information for IS-IS and OSPF, configure traceoptions (see “Tracking Graceful Restart Events” on page 108).

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

```

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

```

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

```

Oct 8 05:20:14 Helper mode for neighbor 192.255.5.1
Oct 8 05:20:14 Received multiple grace lsa from 192.255.5.1

```

Verifying CCC and TCC Graceful Restart

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


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

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

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


CCC Graceful restart : Restarting

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

Example: Configuring Graceful Restart

Figure 7 on page 115 shows a standard MPLS VPN network. Routers CE1 and CE2 are customer edge routers, PE1 and PE2 are provider edge routers, and P0 is a provider core router. Several Layer 3 VPNs are configured across this network, as well as one Layer 2 VPN. Interfaces are shown in the diagram and are not included in the configuration example that follows.

Figure 7: Layer 3 VPN Graceful Restart Topology



Router CE1 On Router CE1, configure the following protocols on the logical interfaces of t3-3/1/0: OSPF on unit 101, RIP on unit 102, BGP on unit 103, and IS-IS on unit 512. Also

configure graceful restart, BGP, IS-IS, OSPF, and RIP on the main instance to be able to connect to the routing instances on Router PE1.

```
[edit]
interfaces {
  t3-3/1/0 {
    encapsulation frame-relay;
    unit 100 {
      dlci 100;
      family inet {
        address 10.96.100.2/30;
      }
    }
    unit 101 {
      dlci 101;
      family inet {
        address 10.96.101.2/30;
      }
    }
    unit 102 {
      dlci 102;
      family inet {
        address 10.96.102.2/30;
      }
    }
    unit 103 {
      dlci 103;
      family inet {
        address 10.96.103.2/30;
      }
    }
    unit 512 {
      dlci 512;
      family inet {
        address 10.96.252.1/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.245.14.172/32;
        primary;
      }
      address 10.96.110.1/32;
      address 10.96.111.1/32;
      address 10.96.112.1/32;
      address 10.96.113.1/32;
      address 10.96.116.1/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4172.00;
    }
  }
}
routing-options {
```

```

    graceful-restart;
    autonomous-system 65100;
}
protocols {
    bgp {
        group CE-PE-INET {
            type external;
            export BGP_INET_LB_DIRECT;
            neighbor 10.96.103.1 {
                local-address 10.96.103.2;
                family inet {
                    unicast;
                }
            }
            peer-as 65103;
        }
    }
}
isis {
    export ISIS_L2VPN_LB_DIRECT;
    interface t3-3/1/0.512;
}
ospf {
    export OSPF_LB_DIRECT;
    area 0.0.0.0 {
        interface t3-3/1/0.101;
    }
}
rip {
    group RIP {
        export RIP_LB_DIRECT;
        neighbor t3-3/1/0.102;
    }
}
}
policy-options {
    policy-statement OSPF_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.101.0/30 exact;
                route-filter 10.96.111.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
    policy-statement RIP_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.102.0/30 exact;
                route-filter 10.96.112.1/32 exact;
            }
            then accept;
        }
    }
}

```

```

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

```

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

```

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

```

```

        address 10.96.101.1/30;
    }
    family mpls;
}
unit 102 {
    dlci 102;
    family inet {
        address 10.96.102.1/30;
    }
    family mpls;
}
unit 103 {
    dlci 103;
    family inet {
        address 10.96.103.1/30;
    }
    family mpls;
}
unit 512 {
    encapsulation frame-relay-ccc;
    dlci 512;
}
}
t1-0/1/0 {
    unit 0 {
        family inet {
            address 10.96.0.2/30;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.245.14.176/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4176.00;
        }
    }
}
}
routing-options {
    graceful-restart;
    router-id 10.245.14.176;
    autonomous-system 69;
}
protocols {
    mpls {
        interface all;
    }
    bgp {
        group PEPE {
            type internal;
            neighbor 10.245.14.182 {
                local-address 10.245.14.176;
            }
        }
    }
}

```

```

        family inet-vpn {
            unicast;
        }
        family l2vpn {
            unicast;
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface t1-0/1/0.0;
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface all;
}
}
policy-options {
    policy-statement STATIC-import {
        from community STATIC;
        then accept;
    }
    policy-statement STATIC-export {
        then {
            community add STATIC;
            accept;
        }
    }
    policy-statement OSPF-import {
        from community OSPF;
        then accept;
    }
    policy-statement OSPF-export {
        then {
            community add OSPF;
            accept;
        }
    }
    policy-statement RIP-import {
        from community RIP;
        then accept;
    }
    policy-statement RIP-export {
        then {
            community add RIP;
            accept;
        }
    }
}
policy-statement BGP-INET-import {

```

```

        from community BGP-INET;
        then accept;
    }
    policy-statement BGP-INET-export {
        then {
            community add BGP-INET;
            accept;
        }
    }
    policy-statement L2VPN-import {
        from community L2VPN;
        then accept;
    }
    policy-statement L2VPN-export {
        then {
            community add L2VPN;
            accept;
        }
    }
    community BGP-INET members target:69:103;
    community L2VPN members target:69:512;
    community OSPF members target:69:101;
    community RIP members target:69:102;
    community STATIC members target:69:100;
}
routing-instances {
    BGP-INET {
        instance-type vrf;
        interface t3-0/0/0.103;
        route-distinguisher 10.245.14.176:103;
        vrf-import BGP-INET-import;
        vrf-export BGP-INET-export;
        routing-options {
            graceful-restart;
            autonomous-system 65103;
        }
        protocols {
            bgp {
                group BGP-INET {
                    type external;
                    export BGP-INET-import;
                    neighbor 10.96.103.2 {
                        local-address 10.96.103.1;
                        family inet {
                            unicast;
                        }
                    }
                    peer-as 65100;
                }
            }
        }
    }
}
L2VPN {
    instance-type l2vpn;
    interface t3-0/0/0.512;
    route-distinguisher 10.245.14.176:512;
}

```

```
vrf-import L2VPN-import;
vrf-export L2VPN-export;
protocols {# There is no graceful-restart statement for Layer 2 VPN instances.
    l2vpn {
        encapsulation-type frame-relay;
        site CE1-ISIS {
            site-identifier 512;
            interface t3-0/0/0.512 {
                remote-site-id 612;
            }
        }
    }
}
OSPF {
    instance-type vrf;
    interface t3-0/0/0.101;
    route-distinguisher 10.245.14.176:101;
    vrf-import OSPF-import;
    vrf-export OSPF-export;
    routing-options {
        graceful-restart;
    }
    protocols {
        ospf {
            export OSPF-import;
            area 0.0.0.0 {
                interface all;
            }
        }
    }
}
RIP {
    instance-type vrf;
    interface t3-0/0/0.102;
    route-distinguisher 10.245.14.176:102;
    vrf-import RIP-import;
    vrf-export RIP-export;
    routing-options {
        graceful-restart;
    }
    protocols {
        rip {
            group RIP {
                export RIP-import;
                neighbor t3-0/0/0.102;
            }
        }
    }
}
STATIC {
    instance-type vrf;
    interface t3-0/0/0.100;
    route-distinguisher 10.245.14.176:100;
    vrf-import STATIC-import;
    vrf-export STATIC-export;
```



```

        routing-options {
            graceful-restart;
            static {
                route 10.96.110.1/32 next-hop t3-0/0/0.100;
            }
        }
    }
}

```

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

```

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

```

```

        disable;
    }
    interface lo0.0 {
        passive;
    }
}
ldp {
    interface all;
}
}

```

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

```

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

```

```

        family inet {
            address 10.96.203.1/30;
        }
        family mpls;
    }
    unit 612 {
        encapsulation frame-relay-ccc;
        dlci 612;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.245.14.182/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4182.00;
        }
    }
}
}
routing-options {
    graceful-restart;
    router-id 10.245.14.182;
    autonomous-system 69;
}
protocols {
    mpls {
        interface all;
    }
    bgp {
        group PEPE {
            type internal;
            neighbor 10.245.14.176 {
                local-address 10.245.14.182;
                family inet-vpn {
                    unicast;
                }
                family l2vpn {
                    unicast;
                }
            }
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface t3-0/0/0.0;
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
}
}
ldp {

```

```

    interface all;
  }
  policy-options {
    policy-statement STATIC-import {
      from community STATIC;
      then accept;
    }
    policy-statement STATIC-export {
      then {
        community add STATIC;
        accept;
      }
    }
    policy-statement OSPF-import {
      from community OSPF;
      then accept;
    }
    policy-statement OSPF-export {
      then {
        community add OSPF;
        accept;
      }
    }
    policy-statement RIP-import {
      from community RIP;
      then accept;
    }
    policy-statement RIP-export {
      then {
        community add RIP;
        accept;
      }
    }
    policy-statement BGP-INET-import {
      from community BGP-INET;
      then accept;
    }
    policy-statement BGP-INET-export {
      then {
        community add BGP-INET;
        accept;
      }
    }
    policy-statement L2VPN-import {
      from community L2VPN;
      then accept;
    }
    policy-statement L2VPN-export {
      then {
        community add L2VPN;
        accept;
      }
    }
    community BGP-INET members target:69:103;
    community L2VPN members target:69:512;
    community OSPF members target:69:101;
  }

```

```

community RIP members target:69:102;
community STATIC members target:69:100;
}
routing-instances {
  BGP-INET {
    instance-type vrf;
    interface t1-0/1/3.203;
    route-distinguisher 10.245.14.182:203;
    vrf-import BGP-INET-import;
    vrf-export BGP-INET-export;
    routing-options {
      graceful-restart;
      autonomous-system 65203;
    }
  }
  protocols {
    bgp {
      group BGP-INET {
        type external;
        export BGP-INET-import;
        neighbor 10.96.203.2 {
          local-address 10.96.203.1;
          family inet {
            unicast;
          }
        }
        peer-as 65200;
      }
    }
  }
}
L2VPN {
  instance-type l2vpn;
  interface t1-0/1/3.612;
  route-distinguisher 10.245.14.182:612;
  vrf-import L2VPN-import;
  vrf-export L2VPN-export;
  protocols {# There is no graceful-restart statement for Layer 2 VPN instances.
    l2vpn {
      encapsulation-type frame-relay;
      site CE2-ISIS {
        site-identifier 612;
        interface t1-0/1/3.612 {
          remote-site-id 512;
        }
      }
    }
  }
}
OSPF {
  instance-type vrf;
  interface t1-0/1/3.201;
  route-distinguisher 10.245.14.182:201;
  vrf-import OSPF-import;
  vrf-export OSPF-export;
  routing-options {
    graceful-restart;
  }
}

```

```

    }
    protocols {
        ospf {
            export OSPF-import;
            area 0.0.0.0 {
                interface all;
            }
        }
    }
}
RIP {
    instance-type vrf;
    interface t1-0/1/3.202;
    route-distinguisher 10.245.14.182:202;
    vrf-import RIP-import;
    vrf-export RIP-export;
    routing-options {
        graceful-restart;
    }
    protocols {
        rip {
            group RIP {
                export RIP-import;
                neighbor t1-0/1/3.202;
            }
        }
    }
}
STATIC {
    instance-type vrf;
    interface t1-0/1/3.200;
    route-distinguisher 10.245.14.182:200;
    vrf-import STATIC-import;
    vrf-export STATIC-export;
    routing-options {
        graceful-restart;
        static {
            route 10.96.210.1/32 next-hop t1-0/1/3.200;
        }
    }
}
}
}

```

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

```

[edit]
interfaces {
    t1-0/0/3 {
        encapsulation frame-relay;
        unit 200 {

```

```

        dlci 200;
        family inet {
            address 10.96.200.2/30;
        }
    }
    unit 201 {
        dlci 201;
        family inet {
            address 10.96.201.2/30;
        }
    }
    unit 202 {
        dlci 202;
        family inet {
            address 10.96.202.2/30;
        }
    }
    unit 203 {
        dlci 203;
        family inet {
            address 10.96.203.2/30;
        }
    }
    unit 512 {
        dlci 512;
        family inet {
            address 10.96.252.2/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.245.14.180/32 {
                primary;
            }
            address 10.96.210.1/32;
            address 10.96.111.1/32;
            address 10.96.212.1/32;
            address 10.96.213.1/32;
            address 10.96.216.1/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4180.00;
        }
    }
}
}
routing-options {
    graceful-restart;
    autonomous-system 65200;
}
protocols {
    bgp {
        group CE-PE-INET {
            type external;

```

```

        export BGP_INET_LB_DIRECT;
        neighbor 10.96.203.1 {
            local-address 10.96.203.2;
            family inet {
                unicast;
            }
            peer-as 65203;
        }
    }
}
isis {
    export ISIS_L2VPN_LB_DIRECT;
    interface t1-0/0/3.612;
}
ospf {
    export OSPF_LB_DIRECT;
    area 0.0.0.0 {
        interface t1-0/0/3.201;
    }
}
rip {
    group RIP {
        export RIP_LB_DIRECT;
        neighbor t1-0/0/3.202;
    }
}
}
policy-options {
    policy-statement OSPF_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.201.0/30 exact;
                route-filter 10.96.211.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
    policy-statement RIP_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.202.0/30 exact;
                route-filter 10.96.212.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
}
policy-statement BGP_INET_LB_DIRECT {
    term direct {

```



```

        from {
            protocol direct;
            route-filter 10.96.203.0/30 exact;
            route-filter 10.96.213.1/32 exact;
        }
        then accept;
    }
    term final {
        then reject;
    }
}
policy-statement ISIS_L2VPN_LB_DIRECT {
    term direct {
        from {
            protocol direct;
            route-filter 10.96.216.1/32 exact;
        }
        then accept;
    }
    term final {
        then reject;
    }
}
}

```

Router PE1 Status Before a Restart

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

```

user@PE1> show bgp neighbor
Peer: 10.96.103.2+3785 AS 65100 Local: 10.96.103.1+179 AS 65103
  Type: External   State: Established   Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.96.110.1      Local ID: 10.96.103.1      Active Holdtime: 90
  Keepalive Interval: 30
  Local Interface: t3-0/0/0.103
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI peer can save forwarding state: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table BGP-INET.inet.0 Bit: 30001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync

```

```

    Active prefixes:          0
    Received prefixes:        0
    Suppressed due to damping: 0
    Last traffic (seconds): Received 8    Sent 3    Checked 3
    Input messages:  Total 15    Updates 0    Refreshes 0    Octets 321
    Output messages: Total 18    Updates 2    Refreshes 0    Octets 450
    Output Queue[2]: 0

Peer: 10.245.14.182+4701 AS 69    Local: 10.245.14.176+179 AS 69
Type: Internal    State: Established    Flags: <>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast l2vpn
Local Address: 10.245.14.176 Holdtime: 90 Preference: 170
Number of flaps: 1
Peer ID: 10.245.14.182    Local ID: 10.245.14.176    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of all end-of-rib markers sent: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete

```

```

RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          0
Received prefixes:        0
Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          0
Received prefixes:        0
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 28   Sent 28   Checked 28
Input messages: Total 2      Updates 0      Refreshes 0      Octets 86
Output messages: Total 13    Updates 10    Refreshes 0      Octets 1073
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

```
user@PE1> show route instance detail
```

```
master:
```

```

Router ID: 10.245.14.176
Type: forwarding          State: Active
Restart State: Complete Path selection timeout: 300
Tables:
inet.0                    : 17 routes (15 active, 0 holddown, 1 hidden)
Restart Complete
inet.3                    : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
iso.0                     : 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
mpls.0                   : 19 routes (19 active, 0 holddown, 0 hidden)
Restart Complete
bgp.l3vpn.0              : 10 routes (10 active, 0 holddown, 0 hidden)
Restart Complete
inet6.0                  : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
bgp.l2vpn.0              : 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```
BGP-INET:
```

```

Router ID: 10.96.103.1
Type: vrf                 State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
t3-0/0/0.103
Route-distinguisher: 10.245.14.176:103
Vrf-import: [ BGP-INET-import ]
Vrf-export: [ BGP-INET-export ]

```

```

Tables:
  BGP-INET.inet.0      : 4 routes (4 active, 0 holddown, 0 hidden)
  Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn           State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.245.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0      : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.96.101.1
  Type: vrf             State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.245.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0        : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.96.102.1
  Type: vrf             State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.245.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0         : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.96.100.1
  Type: vrf             State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.245.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0      : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
__juniper_private1__:
  Router ID: 0.0.0.0
  Type: forwarding      State: Active

user@PE1> show route protocol l2vpn
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
Restart Complete
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)

```

```

Restart Complete
OSPF.inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
mpls.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
800003          *[L2VPN/7] 00:06:00
                 > via t3-0/0/0.512, Pop          Offset: 4
t3-0/0/0.512    *[L2VPN/7] 00:06:00
                 > via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4
bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Complete
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.245.14.176:512:512:611/96
                 *[L2VPN/7] 00:06:01
                 Discard

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

```

Router PE1 Status During a Restart

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

```

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

```

The following sample output is captured during the router restart:

```

user@PE1> show bgp neighbor
Peer: 10.96.103.2      AS 65100 Local: 10.96.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle     Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.245.14.182+179 AS 69      Local: 10.245.14.176+2131 AS 69
  Type: Internal      State: Established  Flags: <ImportEval>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.245.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0

```

```

Peer ID: 10.245.14.182    Local ID: 10.245.14.176    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast 12vpn
NLRI advertised by peer: inet-vpn-unicast 12vpn
NLRI for this session: inet-vpn-unicast 12vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast 12vpn
NLRI peer can save forwarding state: inet-vpn-unicast 12vpn
NLRI that peer saved forwarding for: inet-vpn-unicast 12vpn
NLRI that restart is negotiated for: inet-vpn-unicast 12vpn
NLRI of received end-of-rib markers: inet-vpn-unicast 12vpn
Table bgp.13vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.12vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table L2VPN.12vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1

```

```

Received prefixes: 1
Suppressed due to damping: 0
Last traffic (seconds): Received 0 Sent 0 Checked 0
Input messages: Total 14 Updates 13 Refreshes 0 Octets 1053
Output messages: Total 3 Updates 0 Refreshes 0 Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

user@PE1> **show route instance detail**

master:

```

Router ID: 10.245.14.176
Type: forwarding State: Active
Restart State: Pending Path selection timeout: 300
Tables:
inet.0 : 17 routes (15 active, 1 holddown, 1 hidden)
Restart Pending: OSPF LDP
inet.3 : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: OSPF LDP
iso.0 : 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
mpls.0 : 23 routes (23 active, 0 holddown, 0 hidden)
Restart Pending: LDP VPN
bgp.l3vpn.0 : 10 routes (10 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN
inet6.0 : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
bgp.l2vpn.0 : 1 routes (1 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

```

BGP-INET:

```

Router ID: 10.96.103.1
Type: vrf State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
t3-0/0/0.103
Route-distinguisher: 10.245.14.176:103
Vrf-import: [ BGP-INET-import ]
Vrf-export: [ BGP-INET-export ]
Tables:
BGP-INET.inet.0 : 6 routes (5 active, 0 holddown, 0 hidden)
Restart Pending: VPN

```

L2VPN:

```

Router ID: 0.0.0.0
Type: l2vpn State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
t3-0/0/0.512
Route-distinguisher: 10.245.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
L2VPN.l2vpn.0 : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: VPN L2VPN

```

OSPF:

```

Router ID: 10.96.101.1

```

```

Type: vrf                      State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.245.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Tables:
  OSPF.inet.0                  : 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN
RIP:
Router ID: 10.96.102.1
Type: vrf                      State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.245.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0                   : 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN
STATIC:
Router ID: 10.96.100.1
Type: vrf                      State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.245.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0                : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN
__juniper_private1__:
Router ID: 0.0.0.0
Type: forwarding              State: Active

```

```
user@PE1> show route instance summary
```

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding	inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		l3vpn.0	0/0/0
		inet6.0	2/0/0
		l2vpn.0	0/0/0
		l2circuit.0	0/0/0
BGP-INET	vrf	BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
L2VPN	l2vpn	L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.l2vpn.0	2/0/0
OSPF	vrf	OSPF.inet.0	7/0/0
		OSPF.iso.0	0/0/0
		OSPF.inet6.0	0/0/0


```

RIP                vrf
                   RIP.inet.0        6/0/0
                   RIP.iso.0         0/0/0
                   RIP.inet6.0       0/0/0
STATIC             vrf
                   STATIC.inet.0     4/0/0
                   STATIC.iso.0      0/0/0
                   STATIC.inet6.0    0/0/0
__juniper_private1__ forwarding
                   __juniper_priva.inet.0 0/0/0
                   __juniper_privat.iso.0 0/0/0
                   __juniper_priv.inet6.0 0/0/0

user@PE1> show route protocol l2vpn

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

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

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

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

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

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

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

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

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

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

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

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

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

```


Chapter 16

Summary of Graceful Restart Configuration Statements

This chapter provides a reference for each of the graceful restart configuration statements. The statements are organized alphabetically.

disable

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols (bgp | isis | ldp | ospf | ospf3 | pim | rip | ripng | rsvp) graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (bgp | ldp | ospf | ospf3 | pim) graceful-restart],
[edit protocols (bgp | esis | isis | ospf | ospf3 | ldp | pim | rip | ripng | rsvp) graceful-restart],
[edit protocols bgp group *group-name* graceful-restart],
[edit protocols bgp group *group-name* neighbor *ip-address* graceful-restart],
[edit routing-instances *routing-instance-name* protocols (bgp | ldp | ospf | ospf3 | pim) graceful-restart],
[edit routing-instances *routing-instance-name* routing-options graceful-restart],
[edit routing-options graceful-restart]

Release Information Statement introduced before JUNOS Release 7.4.

Description Disable graceful restart.

Usage Guidelines See “Graceful Restart Configuration Guidelines” on page 103.

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

graceful-restart

Syntax graceful-restart {
 disable;
 helper-disable;
 maximum-helper-recovery-time *seconds*;
 maximum-helper-restart-time *seconds*;
 notify-duration *seconds*;
 recovery-time *seconds*;
 restart-duration *seconds*;
 stale-routes-time *seconds*;
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols (bgp | isis | ldp | ospf | ospf3 | pim | rip | ripng | rsvp)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (bgp | ldp | ospf | ospf3 | pim)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
 [edit protocols (bgp | esis | isis | ldp | ospf | ospf3 | pim | rip | ripng | rsvp)],
 [edit protocols bgp group *group-name*],
 [edit protocols bgp group *group-name* neighbor *ip-address*],
 [edit routing-instances *routing-instance-name* protocols (bgp | ldp | ospf | ospf3 | pim)],
 [edit routing-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable graceful restart.

Options The statements are explained separately.

Usage Guidelines See “Graceful Restart Configuration Guidelines” on page 103.

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

helper-disable

Syntax	helper-disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (isis ldp ospf ospf3 rsvp) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart], [edit protocols (isis ldp ospf ospf3 rsvp) graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable helper mode for graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart.
Default	Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.
Usage Guidelines	See “Graceful Restart Configuration Guidelines” on page 103.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

maximum-helper-recovery-time

Syntax	maximum-helper-recovery-time <i>seconds</i> ;
Hierarchy Level	[edit protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the amount of time the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.
Options	<i>seconds</i> —Amount of time, the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart. Range: 1 through 3600 Default: 180
Usage Guidelines	See “Configuring Graceful Restart Options for RSVP, CCC, and TCC” on page 109.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	maximum-helper-restart-time

maximum-helper-restart-time

Syntax	maximum-helper-restart-time <i>seconds</i> ;
Hierarchy Level	[edit protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Specify the amount of time the router waits after it discovers that a neighboring router has gone down before it declares the neighbor down. This value is applied to all RSVP neighbor routers and should be based on the time that the slowest RSVP neighbor requires for restart.
Options	<i>seconds</i> —The time the router waits after it discovers that a neighboring router has gone down before it declares the neighbor down. Range: 1 through 1800 Default: 60
Usage Guidelines	See “Configuring Graceful Restart Options for RSVP, CCC, and TCC” on page 109.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Topics	maximum-helper-recovery-time

maximum-neighbor-reconnect-time

Syntax	maximum-neighbor-reconnect-time <i>seconds</i> ;
Hierarchy Level	[edit protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify the maximum amount of time allowed to reestablish connection from a restarting neighbor.
Options	<i>seconds</i> —Maximum time allowed for reconnection. Range: 30 through 300
Usage Guidelines	See “Configuring Graceful Restart Options for LDP” on page 110
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.

maximum-neighbor-recovery-time

Syntax	maximum-neighbor-recovery-time <i>seconds</i> ;
Hierarchy Level	[edit protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit routing-instances <i>instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in JUNOS Release 8.3. Statement name changed from <i>maximum-recovery-time</i> to <i>maximum-neighbor-recovery-time</i> in JUNOS Release 9.1.
Description	Specify the amount of time the router retains the state of its Label Distribution Protocol (LDP) neighbors while they undergo a graceful restart.
Options	<i>seconds</i> —Time, in seconds, the router retains the state of its LDP neighbors while they undergo a graceful restart. Range: 140 through 1900 Default: 240
Usage Guidelines	See “Configuring Graceful Restart Options for LDP” on page 110.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Topics	no-strict-lsa-checking recovery-time

no-strict-lsa-checking

Syntax	no-strict-lsa-checking;
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router.
Default	By default, LSA checking is enabled.
Usage Guidelines	See “Configuring Graceful Restart Options for OSPF and OSPFv3” on page 106.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Topics	maximum-neighbor-recovery-time recovery-time

notify-duration

Syntax	notify-duration <i>seconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart], [edit routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Specify the length of time the router notifies helper OSPF routers that it has completed graceful restart.
Options	<i>seconds</i> —Amount of time in the router notifies helper OSPF routers that it has completed graceful restart. Range: 1 through 3600 Default: 30
Usage Guidelines	See “Configuring Graceful Restart Options for OSPF and OSPFv3” on page 106.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Topics	restart-duration

reconnect-time

Syntax	reconnect-time <i>seconds</i> ;
Hierarchy Level	[edit protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Specify the amount of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
Options	<i>seconds</i> —Time required for reconnection. Range: 30 through 300
Usage Guidelines	See “Configuring Graceful Restart Options for LDP” on page 110
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.

recovery-time

Syntax	<code>recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the amount of time a router waits for Label Distribution Protocol (LDP) neighbors to assist it with a graceful restart.
Options	<i>seconds</i> —Time the router waits for LDP to restart gracefully. Range: 120 through 1800 Default: 160
Usage Guidelines	See “Configuring Graceful Restart Options for LDP” on page 110.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	maximum-neighbor-recovery-time no-strict-lsa-checking

restart-duration

Syntax	<code>restart-duration seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (isis ospf ospf3 pim) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</p> <p>[edit protocols (esis isis ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-options graceful-restart]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the duration of the graceful restart period globally. 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.
Options	<p><i>seconds</i>—Time for the graceful restart period.</p> <p>Range: The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none"> ■ [edit routing-options graceful-restart] (global setting)—120 through 900 ■ ES-IS—30 through 300 ■ IS-IS—30 through 300 ■ OSPF/OSPFv3—1 through 3600 ■ PIM—30 through 300 <p>Default: The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none"> ■ [edit routing-options graceful-restart] (global setting)—300 ■ ES-IS—180 ■ IS-IS—210 ■ OSPF/OSPFv3—180 ■ PIM—60
Usage Guidelines	See “Graceful Restart Configuration Guidelines” on page 103.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

restart-time

Syntax	restart-time <i>seconds</i> ;
Hierarchy Level	[edit protocols (bgp rip ripng) graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols (bgp rip ripng) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Configure the duration of the Border Gateway Protocol (BGP), Routing Information Protocol (RIP), or next-generation RIP (RIPng) graceful restart period.
Options	<i>seconds</i> —Amount of time for the graceful restart period. Range: 1 through 600 Default: The range of values varies according the protocol: <ul style="list-style-type: none"> ■ BGP—120 ■ RIP/RIPng—60
Usage Guidelines	See “Configuring Graceful Restart Options for BGP” on page 104 and “Configuring Graceful Restart Options for RIP and RIPng” on page 107.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Topics	stale-routes-time

stale-routes-time

Syntax	stale-routes-time <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-routing-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Configure amount of time the router waits to receive restart messages from restarting Border Gateway Protocol (BGP) neighbors before declaring them down.
Options	<i>seconds</i> —Time the router waits to receive messages from restarting neighbors before declaring them down. Range: 1 through 600 Default: 300
Usage Guidelines	See “Configuring Graceful Restart Options for BGP” on page 104.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Topics	restart-time

traceoptions

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

Default: 128 KB

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

Usage Guidelines See “Tracking Graceful Restart Events” on page 108.

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

Part 7

Virtual Router Redundancy Protocol

- VRRP Overview on page 155
- VRRP Configuration Guidelines on page 157
- Summary of VRRP Configuration Statements on page 175

Chapter 17

VRRP Overview

This chapter contains the following section:

- Understanding VRRP on page 155

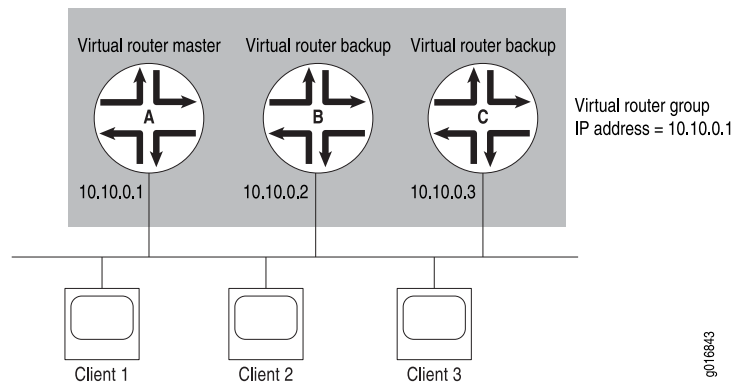
Understanding VRRP

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

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

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

Figure 8 on page 156 illustrates a basic VRRP topology. In this example, Routers A, B, and C are running VRRP and together they make up a virtual router. The IP address of this virtual router is 10.10.0.1 (the same address as the physical interface of Router A).

Figure 8: Basic VRRP

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

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

Chapter 18

VRRP Configuration Guidelines

This chapter contains the following topics:

- VRRP Configuration Hierarchy on page 157
- VRRP for IPv6 Configuration Hierarchy on page 158
- Configuring the Startup Period for VRRP Operations on page 159
- Configuring Basic VRRP Support on page 159
- Configuring VRRP Authentication (IPv4 Only) on page 161
- Configuring the Advertisement Interval for the VRRP Master Router on page 162
- Configuring a Backup Router to Preempt the Master Router on page 163
- Modifying the Preemption Hold-Time Value on page 164
- Configuring an Interface to Accept Packets Destined for the Virtual IP Address on page 164
- Configuring a Logical Interface to Be Tracked on page 165
- Configuring a Route to Be Tracked on page 167
- Configuring Inheritance for a VRRP Group on page 167
- Tracing VRRP Operations on page 168
- Configuring the Silent Period on page 169
- Configuring Passive ARP Learning for Backup VRRP Routers on page 169
- Enabling the Distributed Periodic Packet Management Process for VRRP on page 170
- Example: Configuring VRRP on page 171
- Example: Configuring VRRP for IPv6 on page 172
- Example: Configuring VRRP Route Tracking on page 173

VRRP Configuration Hierarchy

To configure VRRP, include the following statements at either of these hierarchy levels:

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

```

vrrp-group group-id {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      priority-cost priority;
      bandwidth-threshold bits-per-second {
        priority-cost priority;
      }
    }
    priority-hold-time seconds;
    route prefix routing-instance instance-name {
      priority-cost priority;
    }
  }
  virtual-address [ addresses ];
}

```

VRRP for IPv6 Configuration Hierarchy

To configure VRRP for IPv6, include the following statements at either of these hierarchy levels:

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

```

vrrp-inet6-group group-id {
  (accept-data | no-accept-data);
  fast-interval milliseconds;
  inet6-advertise-interval seconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      priority-cost priority;
      bandwidth-threshold bits-per-second {
        priority-cost priority;
      }
    }
    priority-hold-time seconds;
    route prefix routing-instance instance-name {
      priority-cost priority;
    }
  }
}

```

```

    virtual-inet6-address [ addresses ];
    virtual-link-local-address ipv6-address;
}

```

Configuring the Startup Period for VRRP Operations

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

```

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

```

Configuring Basic VRRP Support

An interface can be a member of one or more VRRP groups. To configure basic VRRP support, configure VRRP groups on interfaces by including the `vrrp-group` statement:

```

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

```

You can include this statement at the following hierarchy levels:

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

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

```

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

```

You can include this statement at the following hierarchy levels:

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

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

For each VRRP group, you must configure the following:

- Group identifier—Assign a value from 0 through 255.
- Address of one or more virtual routers that are members of the VRRP group—Normally, you configure only one virtual IP address per group. However,

you can configure up to eight addresses. Do not include a prefix length in a virtual IP address. The following considerations apply to configuring a virtual IP address:

- The virtual router IP address must be the same for all routing platforms in the VRRP group.
- If you configure a virtual IP address to be the same as the physical interface's address, the interface becomes the master virtual router for the group. In this case, you must configure the priority to be 255 and you must configure preemption by including the **preempt** statement.
- If the virtual IP address you choose is not the same as the physical interface's address, you must ensure that the virtual IP address does not appear anywhere else in the routing platform's configuration. Verify that you do not use this address for other interfaces, for the IP address of a tunnel, or for the IP address of static ARP entries.
- You cannot configure a virtual IP address to be the same as the interface's address for an aggregated Ethernet interface. This configuration is not supported.
- For VRRP for IPv6, the EUI-64 option cannot be used. In addition, the Duplicate Address Detection (DAD) process will not run for virtual IPv6 addresses.
- Virtual link local address—(VRRP for IPv6 only) You must explicitly define a virtual link local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link local address must be on the same subnet as the physical interface address.
- Priority for this routing platform to become the master virtual router—Configure the value used to elect the master virtual router in the VRRP group. It can be a number from 1 through 255. The default value for backup routers is 100. A larger value indicates a higher priority. The routing platform with the highest priority within the group becomes the master router.



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



NOTE: If you enable MAC source address filtering on an interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement at the [edit interfaces *interface-name*] hierarchy. (For more information, see the *JUNOS Network Interfaces Configuration Guide*.) MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2378. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Configuring VRRP Authentication (IPv4 Only)

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

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

To enable authentication and specify an authentication method, include the `authentication-type` statement at either of the following hierarchy levels:

```
authentication-type authentication;
```

authentication can be `simple` or `md5`. The authentication type must be the same for all routing platforms in the VRRP group.

You can include this statement at the following hierarchy levels:

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

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

```
authentication-key key;
```

key (the password) is an ASCII string. For simple authentication, it can be from 1 through 8 characters long. For MD5 authentication, it can be from 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (“ ”). The key must be the same for all routing platforms in the VRRP group.

You can include this statement at the following hierarchy levels:

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

Configuring the Advertisement Interval for the VRRP Master Router

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

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

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



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

This topic contains the following sections:

- Modifying the Advertisement Interval in Seconds on page 162
- Modifying the Advertisement Interval in Milliseconds on page 163

Modifying the Advertisement Interval in Seconds

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

```
advertise-interval seconds;
```

The interval can be from 1 through 255 seconds.

You can include this statement at the following hierarchy levels:

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

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

Modifying the Advertisement Interval in Milliseconds

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

```
fast-interval milliseconds;
```

The interval can be from 100 through 999 milliseconds.

You can include this statement at the following hierarchy levels:

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



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

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

```
inet6-advertise-interval ms;
```

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

You can include this statement at the following hierarchy levels:

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

Configuring a Backup Router to Preempt the Master Router

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

```
preempt;
```

You can include this statement at the following hierarchy levels:

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

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

```
no-preempt;
```

Modifying the Preemption Hold-Time Value

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

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

To modify the preemption hold-time value, include the **hold-time** statement at either of the following hierarchy levels:

```
hold-time seconds;
```

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy levels:

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

Configuring an Interface to Accept Packets Destined for the Virtual IP Address

To configure an interface to accept packets destined for the virtual IP address, include the **accept-data** statement:

```
accept-data;
```

You can include this statement at the following hierarchy levels:

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

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

To prohibit the interface from accepting packets destined for the virtual IP address, include the **no-accept-data** statement:

```
no-accept-data;
```

Including the **accept-data** statement has the following consequences:

- If the master router owns the virtual IP address, the **accept-data** statement is not valid.
- If the priority of the master router is set to 255, the **accept-data** statement is not valid.
- To restrict incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.
- If the master router owns the virtual IP address, the master router responds to Internet Control Message Protocol (ICMP) message requests only.
- If you include the **accept-data** statement:
 - Your routing platform configuration does not comply with RFC 2378.
 - VRRP clients can process gratuitous ARP.
 - VRRP clients must not use packets other than ARP replies to update their ARP cache.

Configuring a Logical Interface to Be Tracked

VRRP can track whether a logical interface is up, down, or not present, and can also dynamically change the priority of the VRRP group based on the state of the tracked logical interface, which might trigger a new master router election. VRRP can also track the operational speed of a logical interface and dynamically update the priority of the VRRP group when the speed crosses a configured threshold.

When interface tracking is enabled, you cannot configure a priority of 255, thereby designating the master router. For each VRRP group, you can track up to 10 logical interfaces.

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

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

You can include these statements at the following hierarchy levels:

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

The interface specified is the interface to be tracked for the VRRP group. The priority hold time is the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.

The bandwidth threshold specifies a threshold for the tracked interface. When the bandwidth of the tracked interface drops below the configured bandwidth threshold value, the VRRP group uses the bandwidth threshold priority cost. You can track up to five bandwidth threshold statements for each tracked interface.

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked logical interface goes down, forcing a new master router election. The value can be from 1 through 254. The sum of the costs for all tracked logical interfaces or routes must be less than or equal to the configured priority of the VRRP group.

If you are tracking more than one interface, the router applies the sum of the priority costs for the tracked interfaces (at most, only one priority cost for each tracked interface) to the VRRP group priority. However, the interface priority cost and bandwidth threshold priority cost values for each VRRP group are not cumulative. The router uses only one priority cost to a tracked interface as indicated in Table 8 on page 166:

Table 8: Interface State and Priority Cost Usage

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

You must configure an interface priority cost only if you have configured no bandwidth thresholds. If you have not configured an interface priority cost value, and the interface is down, the interface uses the bandwidth threshold priority cost value of the lowest bandwidth threshold.

Configuring a Route to Be Tracked

VRRP can track whether a route is reachable (that is, the route exists in the routing table of the routing instance included in the configuration) and dynamically change the priority of the VRRP group based on the reachability of the tracked route, which might trigger a new master router election.

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

```
track {
  priority-hold-time seconds;
  route prefix routing-instance instance-name {
    priority-cost priority;
  }
}
```

You can include these statements at the following hierarchy levels:

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

The route prefix specified is the route to be tracked for the VRRP group. The priority hold time is the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.

The routing instance is the routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, specify the instance name as **default**.



NOTE: Tracking a route that belongs to a routing instance from a different logical system is not supported.

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked route goes down, forcing a new master router election. The value can be from 1 through 254. The sum of the costs for all tracked logical interfaces or routes must be less than or equal to the configured priority of the VRRP group.

Configuring Inheritance for a VRRP Group

JUNOS Software enables you to configure VRRP groups on the various subnets of a VLAN to inherit the state and configuration of one of the groups, which is known as the *active VRRP group*. By configuring inheritance, you can prevent VRRP groups other than the active group from sending out VRRP advertisements. When the

vrp-inherit-from configuration statement is included in the configuration, only the active VRRP group from which the other VRRP groups are inheriting the state sends out VRRP advertisements; the groups inheriting the state do not send any VRRP advertisements, because the state is maintained only on the group from which the state is inherited.

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

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

```
vrp-inherit-from vrrp-group
```

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

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

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

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

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

Tracing VRRP Operations

To trace VRRP operations, include the **traceoptions** statement at the [edit protocols *vrp*] hierarchy level.

By default, VRRP logs the error, data carrier detect (DCD) configuration, and routing socket events in a file in the `/var/log` directory. By default, this file is named `/var/log/vrrpd`. The default file size is 1 megabyte (MB), and three files are created before the first one gets overwritten.

To change the configuration of the logging file, include the `traceoptions` statement at the `[edit protocols vrrp]` hierarchy level:

```
[edit protocols vrrp]
traceoptions {
  file <filename> <files number> <match regular-expression> <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
flag flag;
```

You can specify the following VRRP tracing flags:

- `all`—Trace all VRRP operations.
- `database`—Trace all database changes.
- `general`—Trace all general events.
- `interfaces`—Trace all interface changes.
- `normal`—Trace all normal events.
- `packets`—Trace all packets sent and received.
- `state`—Trace all state transitions.
- `timer`—Trace all timer events.

Configuring the Silent Period

The silent period starts when the interface state is changed from down to up. During this period, the Master Down Event is ignored. Configure the silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets during the interface startup phase.

To configure the silent period interval that the Master Down Event timer ignores, include the `startup-silent-period` statement at the `[edit protocols vrrp]` hierarchy level:

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

Configuring Passive ARP Learning for Backup VRRP Routers

By default, the backup VRRP router drops ARP requests for the VRRP-IP to VRRP-MAC address translation. This means that the backup router does not learn the ARP (IP-to-MAC address) mappings for the hosts sending the requests. When it detects a failure of the master router and transitions to become the new master router, the backup router must learn all the entries that were present in the ARP cache of the

master router. In environments with many directly attached hosts, such as metro Ethernet environments, the number of ARP entries to learn can be high. This can cause a significant transition delay, during which the traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup router to hold approximately the same contents as the ARP cache in the master router, thus preventing the problem of learning ARP entries in a burst. To enable passive ARP learning, include the `passive-learning` statement at the `[edit system arp]` hierarchy level:

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

We recommend setting passive learning on both the backup and master VRRP routers. Doing so prevents the need to manually intervene when the master router becomes the backup router. While a router is operating as the master router, the passive learning configuration has no operational impact. The configuration takes effect only when the router is operating as a backup router.

For information about configuring gratuitous ARP and the ARP aging timer, see the *JUNOS System Basics Configuration Guide*.

Enabling the Distributed Periodic Packet Management Process for VRRP

Typically, VRRP advertisements are sent by the VRRP process (`vrrpd`) on the master VRRP router at regular intervals to let other members of the group know that the VRRP master router is operational.

When the `vrrpd` process is busy and does not send VRRP advertisements, the backup VRRP routers may assume that the master router is down and take over as the master router, causing unnecessary flaps. This takeover may occur even though the original master router is still active and available, and might resume sending advertisements after the traffic has decreased. To address this problem and to reduce the load on the `vrrpd` process, JUNOS Software now uses the periodic packet management process (`ppmd`) to send VRRP advertisements on behalf of the `vrrpd` process. However, you can further delegate the job of sending VRRP advertisements to the distributed `ppmd` process that resides on the Packet Forwarding Engine.

The ability to delegate the sending of VRRP advertisements to the distributed `ppmd` process ensures that the VRRP advertisements are sent even when the `ppmd` process—which is now responsible for sending VRRP advertisements—is busy. Such delegation prevents the possibility of false alarms when the `ppmd` process is busy. The ability to delegate the sending of VRRP advertisements to distributed `ppmd` also adds to scalability because the load is shared across multiple `ppmd` instances and is not concentrated on any single unit.

To configure the distributed `ppmd` process to send VRRP advertisements when the `ppmd` process is busy, include the `delegate-processing` statement at the `[edit protocols vrrp]` hierarchy level.

```
[edit protocols vrrp]
delegate-processing;
```


Example: Configuring VRRP

Configure one master (Router A) and one backup (Router B) routing platform. The address configured in the **virtual-address** statements differs from the addresses configured in the **address** statements. When you configure multiple VRRP groups on an interface, you configure one to be the master virtual router for that group.

On Router A

```
[edit]
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.20/24 {
          vrrp-group 27 {
            virtual-address 192.168.1.15;
            priority 254;
            authentication-type simple;
            authentication-key booJUM;
          }
        }
      }
    }
  }
}
```

On Router B

```
[edit]
interfaces {
  ge-4/2/0 {
    unit 0 {
      family inet {
        address 192.168.1.24/24 {
          vrrp-group 27 {
            virtual-address 192.168.1.15;
            priority 200;
            authentication-type simple;
            authentication-key booJUM;
          }
        }
      }
    }
  }
}
```

Configuring One Router to Be the Master Virtual Router for the Group

```
[edit]
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.20/24 {
          vrrp-group 2 {
            virtual-address 192.168.1.20;
            priority 255;
            advertise-interval 3;
          }
        }
      }
    }
  }
}
```

```
    preempt;
}
vrrp-group 10 {
    virtual-address 192.168.1.55;
    priority 201;
    advertise-interval 3;
}
vrrp-group 1 {
    virtual-address 192.168.1.54;
    priority 22;
    advertise-interval 4;
}
}
}
}
```

Configuring VRRP and MAC Source Address Filtering

The VRRP group number is the decimal equivalent of the last byte of the virtual MAC address.

```
[edit interfaces]
ge-5/2/0 {
  gigether-options {
    source-filtering;
    source-address-filter {
      00:00:5e:00:01:0a; # Virtual MAC address
    }
  }
}
unit 0 {
  family inet {
    address 192.168.1.10/24 {
      vrrp-group 10 { # VRRP group number
        virtual-address 192.168.1.10;
        priority 255;
        preempt;
      }
    }
  }
}
```

Example: Configuring VRRP for IPv6

Configure VRRP properties for IPv6 in one master (Router A) and one backup (Router B).

On Router A

```
[edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet6 {
      address fe80::5:0:0:6/64;
      address fec0::5:0:0:6/64 {
        vrp-inet6-group 3 { # VRRP inet6 group number
          virtual-inet6-address fec0::5:0:0:7;
```

```

        virtual-link-local-address fe80::5:0:0:7;
        priority 200;
        preempt;
    }
}

[edit protocols]
router-advertisement {
    interface ge-1/0/0.0 {
        prefix fec0::/64;
        max-advertisement-interval 4;
    }
}

On Router B [edit interfaces]
ge-1/0/0 {
    unit 0 {
        family inet6 {
            address fe80::5:0:0:8/64;
            address fec0::5:0:0:8/64 {
                vrrp-inet6-group 3 { # VRRP inet6 group number
                    virtual-inet6-address fec0::5:0:0:7;
                    virtual-link-local-address fe80::5:0:0:7;
                    priority 100;
                    preempt;
                }
            }
        }
    }
}

[edit protocols]
router-advertisement {
    interface ge-1/0/0.0 {
        prefix fec0::/64;
        max-advertisement-interval 4;
    }
}

```

Example: Configuring VRRP Route Tracking

Configure routers R1 and R2 to run VRRP. Configure static routes and a policy for exporting the static routes on R3. The VRRP routing instances on R2 track the routes that are advertised by R3.

R1 Configuration

```

[edit interfaces]
ge-1/0/3 {
    unit 0 {
        vlan-id1;
        family inet {
            address 200.100.50.2/24 {
                vrrp-group 0 {
                    virtual-address 200.100.50.101;

```

```

        priority 195;
    }
}
}
}
}

```

R2 Configuration

```

[edit interfaces]
ge-1/0/1 {
  unit 0 {
    vlan-id 1;
    family inet {
      address 200.100.50.1/24 {
        vrrp-group 0 {
          virtual-address 200.100.50.101;
          priority 200;
          track {
            route 59.0.58.153/22 routing-instance default priority-cost 5
            route 59.0.58.154/32 routing-instance default priority-cost 5
            route 59.0.58.155/32 routing-instance default priority-cost 5
          }
        }
      }
    }
  }
}
}
}
}

```

R3 Configuration

```

[edit]
policy-options {
  policy-statement static-policy {
    term term1 {
      then accept;
    }
  }
}
protocols {
  ospf {
    export static-policy;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
routing-options {
  static {
    route 59.0.0.153/32 next-hop 45.45.45.46;
    route 59.0.0.154/32 next-hop 45.45.45.46;
    route 59.0.0.155/32 next-hop 45.45.45.46;
  }
}
}

```

Chapter 19

Summary of VRRP Configuration Statements

This chapter provides a reference for each of the VRRP configuration statements. The statements are organized alphabetically.

accept-data

Syntax	(accept-data no-accept-data);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not an interface accepts packets destined for the virtual IP address:</p> <ul style="list-style-type: none">■ accept-data—Enable the interface to accept packets destined for the virtual IP address.■ no-accept-data—Prevent the interface from accepting packets destined for the virtual IP address.
Default	If the accept-data statement is not configured, the master router responds to Internet Control Message Protocol (ICMP) message requests only.
Usage Guidelines	See “Configuring an Interface to Accept Packets Destined for the Virtual IP Address” on page 164.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

advertise-interval

Syntax	advertise-interval <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets. All routers in the VRRP group must use the same advertisement interval.
Options	<i>seconds</i> —Interval between advertisement packets. Range: 1 through 255 seconds Default: 1 second
Usage Guidelines	See “Configuring the Advertisement Interval for the VRRP Master Router” on page 162.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	fast-interval inet6-advertise-interval

authentication-key

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement.</p> <p>All routers in the VRRP group must use the same authentication scheme and password.</p>
Options	<i>key</i> —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 (" ").
Usage Guidelines	See “Configuring VRRP Authentication (IPv4 Only)” on page 161.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	authentication-type

authentication-type

Syntax	<code>authentication-type authentication;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the authentication-key statement.</p> <p>All routers in the VRRP group must use the same authentication scheme and password.</p>
Options	<p>authentication—Authentication scheme:</p> <ul style="list-style-type: none"> ■ simple—Use a simple password. The password is included in the transmitted packet, making this method of authentication relatively insecure. ■ md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme. <p>Default: none (No authentication is performed)</p>
Usage Guidelines	See “Configuring VRRP Authentication (IPv4 Only)” on page 161.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	authentication-key

bandwidth-threshold

Syntax	<code>bandwidth-threshold <i>bits-per-second</i> { priority-cost <i>priority</i>; }</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i> track interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i> track interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
Options	<i>bits-per-second</i> —Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track. Range: 1 through 10000000000000 bits per second
	The remaining statement is described separately.
Usage Guidelines	See “Configuring a Logical Interface to Be Tracked” on page 165.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	priority-cost

fast-interval

Syntax	<code>fast-interval milliseconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets. All routers in the VRRP group must use the same advertisement interval.
Options	<i>milliseconds</i> —Interval between advertisement packets. Range: 100 through 999 milliseconds Default: 1 second
Usage Guidelines	See “Configuring the Advertisement Interval for the VRRP Master Router” on page 162.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	advertise-interval inet6-advertise-interval

hold-time

Syntax	hold-time <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp vrrp-inet6-group) <i>group-id</i> preempt], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp vrrp-inet6-group) <i>group-id</i> preempt]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.
Default	VRRP preemption is not timed.
Options	<i>seconds</i> —Hold-time period. Range: 0 through 3600 seconds Default: 0 seconds (VRRP preemption is not timed.)
Usage Guidelines	See “Configuring a Backup Router to Preempt the Master Router” on page 163.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

inet6-advertise-interval

Syntax	inet6-advertise-interval <i>ms</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in JUNOS Release 8.4R2.
Description	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets. All routers in the VRRP group must use the same advertisement interval.
Options	<i>ms</i> —Interval, in milliseconds, between advertisement packets. Range: 100 to 40,950 milliseconds (ms) Default: 1 second
Usage Guidelines	See “Configuring the Advertisement Interval for the VRRP Master Router” on page 162.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	fast-interval advertise-interval

interface

Syntax	<pre>interface <i>interface-name</i> { priority <i>priority</i>; bandwidth-threshold <i>bits-per-second</i> { priority-cost <i>priority</i>; } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i> track]
Release Information	Statement introduced in JUNOS Release 7.5. bandwidth-threshold statement added in JUNOS Release 8.1.
Description	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	interface <i>interface-name</i> —Interface to be tracked for this VRRP group. Range: 1 through 10 interfaces The remaining statements are described separately.
Usage Guidelines	See “Configuring a Logical Interface to Be Tracked” on page 165.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

no-accept-data

See accept-data

no-preempt

See preempt

preempt

Syntax	(preempt no-preempt) { hold-time <i>seconds</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a master router:</p> <ul style="list-style-type: none"> ■ preempt—Allow the master router to be preempted. ■ no-preempt—Prohibit the preemption of the master router. <p>The remaining statement is explained separately.</p>
Default	If you omit this statement, the backup router cannot preempt a master router.
Usage Guidelines	See “Configuring a Backup Router to Preempt the Master Router” on page 163.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

priority

Syntax	<code>priority <i>priority</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	<i>priority</i> —Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected. Range: 1 through 255 Default: 100 (for backup routers)
Usage Guidelines	See "Configuring Basic VRRP Support" on page 159.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

priority-cost

Syntax	<code>priority-cost priority;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp vrrp-inet-group) <i>group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp vrrp-inet6-group) <i>group-id</i> track interface <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp vrrp-inet6-group) <i>group-id</i> track interface <i>interface-name</i> bandwidth-threshold <i>bits-per-second</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp vrrp-inet6-group) <i>group-id</i> track interface <i>interface-name</i> bandwidth-threshold <i>bits-per-second</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i> track route <i>prefix</i> routing-instance <i>instance-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 inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i> track route <i>prefix</i> routing-instance <i>instance-name</i>]</p>
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	<p><i>priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p> <p>Range: 1 through 254</p>
Usage Guidelines	See "Configuring a Logical Interface to Be Tracked" on page 165.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

priority-hold-time

Syntax	<code>priority-hold-time seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp vrrp-inet6-group) <i>group-id</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp vrrp-inet6-group) <i>group-id</i> track]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	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. 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.
Options	<i>seconds</i> —The minimum length of time that must elapse between dynamic priority changes. Range: 1 through 3600 seconds
Usage Guidelines	See “Configuring a Logical Interface to Be Tracked” on page 165.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

route

Syntax	route <i>prefix</i> routing-instance <i>instance-name</i> { priority-cost <i>priority</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i> track]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p><i>prefix</i>—Route to be tracked for this VRRP group.</p> <p>routing-instance <i>instance-name</i>—Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for <i>instance-name</i> must be default.</p> <p>The remaining statement is described separately.</p>
Usage Guidelines	See “Configuring a Route to Be Tracked” on page 167.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.

startup-silent-period

Syntax	startup-silent-period <i>seconds</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Instruct the system to ignore the Master Down Event when an interface transitions from the disabled state to the enabled state. This statement is used to avoid an incorrect error alarm caused by delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase.
Options	<p><i>seconds</i>—Number of seconds.</p> <p>Default: 4 seconds</p> <p>Range: 1 through 2000 seconds</p>
Usage Guidelines	See “Configuring the Startup Period for VRRP Operations” on page 159.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

traceoptions

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

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced before JUNOS Release 7.4.

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

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

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

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

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

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.

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

Default: 3 files

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

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

- all—All VRRP tracing operations
- database—Database changes
- general—General events
- interfaces—Interface changes
- normal—Normal events
- packets—Packets sent and received

- **state**—State transitions
- **timer**—Timer events

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

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

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 | no-world-readable—Specifies whether any reader can read the log file.

Usage Guidelines See “Tracing VRRP Operations” on page 168.

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

track

Syntax	<pre>track { interface <i>interface-name</i> { priority-cost <i>priority</i>; bandwidth-threshold <i>bits-per-second</i> { priority-cost <i>priority</i>; } } priority-hold-time <i>seconds</i>; route <i>prefix</i> routing-instance <i>instance-name</i> { priority-cost <i>priority</i>; } }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <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 inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-id</i>]</p>
Release Information	<p>Statement introduced before JUNOS Release 7.4.</p> <p>bandwidth-threshold option added in JUNOS Release 8.1.</p> <p>route option added in JUNOS Release 9.0.</p>
Description	Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.
Options	The remaining statements are described separately.
Usage Guidelines	See “Configuring a Logical Interface to Be Tracked” on page 165 and “Configuring a Route to Be Tracked” on page 167.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

virtual-address

Syntax	virtual-address [<i>addresses</i>];
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 group. You can configure up to eight addresses.
Options	<i>addresses</i> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.
Usage Guidelines	See “Configuring Basic VRRP Support” on page 159.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

virtual-inet6-address

Syntax	virtual-inet6-address [<i>addresses</i>];
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.
Options	<i>addresses</i> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.
Usage Guidelines	See “Configuring Basic VRRP Support” on page 159.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

virtual-link-local-address

Syntax	<code>virtual-link-local-address <i>ipv6-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure a virtual link local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link local address for each VRRP for IPv6 group. The virtual link local address must be in the same subnet as the physical interface address.
Options	<i>ipv6-address</i> —Virtual link local IPv6 address for VRRP for an IPv6 group. Range: 0 through 255 The remaining statements are explained separately.
Usage Guidelines	See “Configuring Basic VRRP Support” on page 159.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vrrp-group

Syntax	<pre> vrrp-group <i>group-id</i> { (accept-data no-accept-data); advertise-interval <i>seconds</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> { priority-cost <i>priority</i>; bandwidth-threshold <i>bits-per-second</i> { priority-cost <i>priority</i>; } } priority-hold-time <i>seconds</i>; route <i>prefix</i> routing-instance <i>instance-name</i> { priority-cost <i>priority</i>; } } virtual-address [<i>addresses</i>]; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>
Release Information	<p>Statement introduced before JUNOS Release 7.4. <i>bandwidth-threshold</i> option added in JUNOS Release 8.1. <i>route</i> option added in JUNOS Release 9.0.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group.</p>
Options	<p><i>group-id</i>—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>Range: 0 through 255</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	<p>See “VRRP Configuration Guidelines” on page 157.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>

vrrp-inet6-group

Syntax	<pre> vrrp-inet6-group <i>group-id</i> { (accept-data no-accept-data); fast-interval <i>milliseconds</i>; inet6-advertise-interval <i>seconds</i>; (preempt no-preempt) { hold-time <i>seconds</i>; } priority <i>number</i>; track { interface <i>interface-name</i> { priority-cost <i>priority</i>; bandwidth-threshold <i>bits-per-second</i> { priority-cost <i>priority</i>; } } priority-hold-time <i>seconds</i>; route <i>prefix</i> routing-instance <i>instance-name</i> { priority-cost <i>priority</i>; } } virtual-inet6-address [<i>addresses</i>]; virtual-link-local-address <i>ipv6-address</i>; } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</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>]
Release Information	Statement introduced before JUNOS Release 7.4. bandwidth-threshold option added in JUNOS Release 8.1. route option added in JUNOS Release 9.0.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group. .
Options	<p><i>group-id</i>—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>Range: 0 through 255</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “VRRP Configuration Guidelines” on page 157.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

Part 8

Unified ISSU

- Unified ISSU Overview on page 199
- Unified ISSU Configuration Guidelines on page 217
- Unified ISSU Configuration Statements Summary on page 235

Chapter 20

Unified ISSU Overview

This chapter includes the following sections:

- Unified ISSU Concepts on page 199
- Unified ISSU Process on the TX Matrix Router on page 204
- Unified ISSU System Requirements on page 205

Unified ISSU Concepts

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

A unified ISSU provides the following benefits:

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



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

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



NOTE: You can verify the unified ISSU-compatibility of the software, hardware, and the configuration on a device by issuing the **request system software validate in-service-upgrade** command. This command runs the validation checks, and shows whether the operating system, device components, and configurations are ISSU compatible or not. For more information about the **request system software validate in-service-upgrade** command, see *JUNOS System Basics and Services Command Reference*.

To perform a unified ISSU, complete the following steps:

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

A JUNOS release package comprises three distinct systems:

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

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

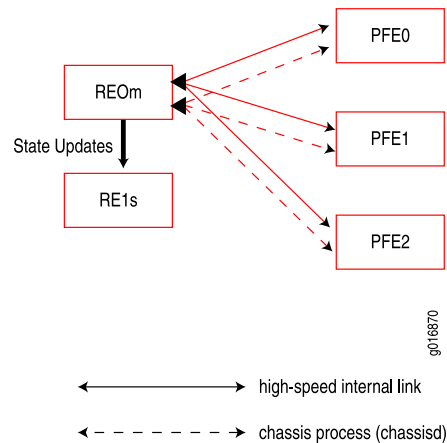


NOTE: In the illustrations, a solid line indicates the high-speed internal link between a Routing Engine and a Packet Forwarding Engine. A dotted line indicates the chassis process (chassisd), another method of communication between a Routing Engine and a Packet Forwarding Engine. RE0m and RE1s indicate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




NOTE: The following process pertains to all supported routing platforms except the TX Matrix router. For information about the unified ISSU process on the TX Matrix router, see “Unified ISSU Process on the TX Matrix Router” on page 204. On M320 and T320 routers and on T640 and T1600 routers, the Packet Forwarding Engine resides on an FPC. However, on an M120 router, the Forwarding Engine Board (FEB) replaces the functions of a Packet Forwarding Engine. In the illustrations and steps, when considering an M120 router, you can regard the PFE as an FPC. As an additional step on an M120 router, after the FPCs and PICs have been upgraded, the FEBs are upgraded.

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

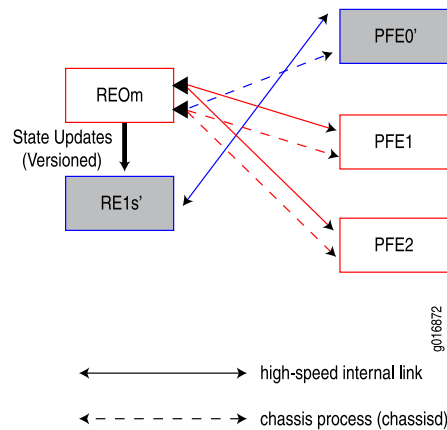


2. When the validation succeeds, the kernel state synchronization daemon (ksyncd) synchronizes the kernel on the backup Routing Engine with the master Routing Engine.
3. The backup Routing Engine is upgraded with the new software image. Before being upgraded, the backup Routing Engine gets the configuration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

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

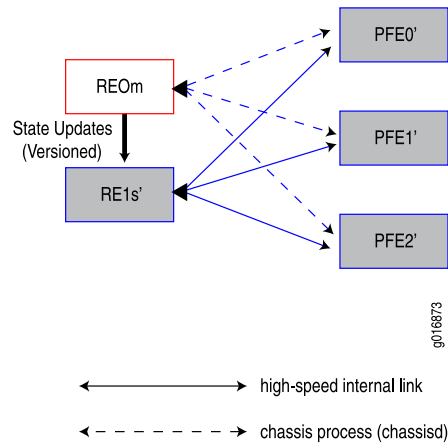


4. The chassis process (chassisd) on the master Routing Engine prepares other software processes for the unified ISSU. When all the processes are ready, chassisd sends an ISSU_PREPARE message to the Flexible PIC Concentrators (FPCs) installed in the router.
5. The Packet Forwarding Engine on each FPC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




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

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




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

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



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



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

Unified ISSU Process on the TX Matrix Router

After you issue the `request system software in-service-upgrade` command on a TX Matrix router, the following process occurs.

1. The management process (mgd) on the master Routing Engine of the TX Matrix router (global master) checks whether nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled in the current configuration.
2. After successful validation of nonstop active routing and graceful Routing Engine switchover configuration, the management process copies the new image to the backup Routing Engines on the TX Matrix router and the T640 routers.
3. The kernel synchronization process (ksyncd) on the backup Routing Engines synchronizes the kernel on the backup Routing Engines with that of the master Routing Engines.
4. The backup Routing Engines are upgraded with the new software and are rebooted. After rebooting, the backup Routing Engines are once again synchronized with the global master Routing Engine.
5. The unified ISSU control moves from the management process to the chassis process (chassisd). The chassis process informs the software processes about the unified ISSU and waits for responses from various software processes (such as spmb).
6. After receiving messages from the software processes indicating that the processes are ready for unified ISSU, the chassis process on the global master Routing Engine sends messages to the chassis process on the routing nodes to start the unified ISSU.
7. The chassis process on the routing nodes sends ISSU_PREPARE messages to the field replaceable units (FRUs), such as FPC and intelligent PICs.

8. On receiving an ISSU_PREPARE message, the Packet Forwarding Engines save the current state information and download the new software image from the backup Routing Engines. Next, each Packet Forwarding Engine sends ISSU_READY messages to the chassis process.
9. On receiving an ISSU_READY message from the Packet Forwarding Engines, the chassis process sends an ISSU_REBOOT message to the FRUs. While the upgrade is in progress, the FRUs keep sending ISSU_IN_PROGRESS messages to the chassis process on the routing nodes. The chassis process on each routing node, in turn, sends an ISSU_IN_PROGRESS message to the chassis process on the global master Routing Engine.
10. After the reboot, the Packet Forwarding Engines restore the saved state information and connect back to the routing nodes; the chassis process on each routing node sends an ISSU_READY message to the chassis process on the global master Routing Engine. The ISSU_READY message from the chassis process on the routing nodes indicates that the unified ISSU is complete on the FRUs.
11. The unified ISSU control moves back to the management process on the global master Routing Engine.
12. The management process initiates Routing Engine switchover on the master Routing Engines.
13. Routing Engine switchover occurs on the TX Matrix router and the T640 routers.



NOTE: Currently, the FRUs on a TX Matrix router do not support graceful Routing Engine switchover and are rebooted every time graceful Routing Engine switchover occurs.

14. After the switchover, the FRUs connect to the new master Routing Engines, and the chassis manager and PFE manager on the T640 router FRUs connect to the new master Routing Engines on the T640 routers.
15. The management process on the global master Routing Engine initiates the upgrade process on the old master Routing Engines on the T640 routers.
16. After the old master Routing Engines on the T640 routers are upgraded, the management process initiates the upgrade of the old global master Routing Engine, that is, the old master Routing Engine on the TX Matrix router.
17. After a successful unified ISSU, the TX Matrix router and the T640 routers are rebooted.

Unified ISSU System Requirements

This section contains the following topics:

- Unified ISSU JUNOS Software Release Support on page 206
- Unified ISSU Platform Support on page 206
- Unified ISSU Protocol Support on page 206
- Unified ISSU Feature Support on page 208

- Unified ISSU PIC Support on page 208
- Unified ISSU DPC and FPC Support on MX Series Routers on page 215

Unified ISSU JUNOS Software Release Support

In order to perform a unified ISSU, your router must be running a JUNOS Software release that supports unified ISSU for the specific platform. See “Unified ISSU Platform Support” on page 206. You can use unified ISSU to upgrade from an ISSU-capable software release to a newer software release. To downgrade from an ISSU-capable release to a previous software release (ISSU-capable or not), use the **request system add** command. See the *JUNOS Software Installation and Upgrade Guide* for details.



NOTE: Unified ISSU does not support extension application packages developed using the Juniper Partner Solution Development Platform (PSDP) SDK.

Unified ISSU Platform Support

Table 9 on page 206 lists the platforms on which a unified ISSU is supported.

Table 9: Unified ISSU Platform Support

Routing Platform	JUNOS Software Release
M120 router	9.2 or later
M320 router	9.0 or later
M10i router with Enhanced Compact Forwarding Engine Board (CFEB-E)	9.5 or later
MX Series Ethernet Services Router	9.3 or later
NOTE: Unified ISSU for MX Series routers does not support IEEE 802.1ag OAM and IEEE 802.3ah protocols.	
T320 router	9.0 or later
T640 router	9.0 or later
T1600 router	9.1 or later
TX Matrix router	9.3 or later

Unified ISSU Protocol Support

Unified ISSU is dependent on nonstop active routing. Table 10 on page 207 lists the protocols that are supported during a unified ISSU.

Table 10: Unified ISSU Protocol Support

Protocol	JUNOS Software Release
BGP, except for BGP VPN services	9.0 or later
IS-IS	9.0 or later
LDP	9.0 or later
LDP-based virtual private LAN service (VPLS)	9.3 or later
Layer 2 circuits	9.2 or later
Layer 3 VPNs using LDP	9.2 or later
Link Aggregation Control Protocol (LACP) on MX Series routers	9.4 or later
OSPF/OSPFv3	9.0 or later
Protocol Independent Multicast (PIM)	9.3 or later

Unified ISSU Support for Layer 2 Control Protocol Process

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

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



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

Unified ISSU enhancements and nonstop active bridging support for the l2cpd process ensure that the new master Routing Engine is able to take control during unified

ISSU without any disruptions in the control plane and minimize the disruptions in the Layer 2 data plane during unified ISSU.

Unified ISSU Feature Support

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

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

Unified ISSU PIC Support

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

- PIC Considerations on page 209
- SONET/SDH PICs on page 209
- Fast Ethernet and Gigabit Ethernet PICs on page 211
- Channelized PICs on page 212
- Tunnel Services PICs on page 213
- ATM PICs on page 213
- Serial PICs on page 214
- DS3, E1, E3, and T1 PICs on page 214
- Enhanced IQ PICs on page 215
- Enhanced IQ2 Ethernet Services Engine (ESE) PIC on page 215



NOTE: For information about FPC types, FPC/PIC compatibility, and the initial JUNOS Software release in which an FPC supported a particular PIC, see the PIC guide for your router platform.

PIC Considerations

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

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

SONET/SDH PICs

Table 11 on page 210 lists the SONET/SDH PICs that are supported during a unified ISSU.

Table 11: Unified ISSU PIC Support: SONET/SDH

PIC Type	Number of Ports	Model Number	Router
OC3c/STM1	4-port	PB-4OC3-SON-MM — (EOL)	M120 M320, T320, T640, T1600
		PB-4OC3-SON-SMIR — (EOL)	
		PE-4OC3-SON-MM — (EOL)	M10i
		PE-4OC3-SON-SMIR — (EOL)	
	2-port	PE-2OC3-SON-MM — (EOL)	
		PE-2OC3-SON-SMIR — (EOL)	
OC3c/STM1 with SFP	2-port	PE-2OC3-SON-SFP	M10i
OC3c/STM1, SFP (Multi-Rate)	4 OC3 ports, 4 OC12 ports	PB-4OC3-4OC12-SON-SFP	M120 M320, T320, T640, T1600
	4 OC3 ports, 1 OC12 port	PB-4OC3-1OC12-SON2-SFP	
		PE-4OC3-1OC12-SON-SFP	M10i
OC12c/STM4	1-port	PE-1OC12-SON-SFP	M10i
		PE-1OC12-SON-MM — (EOL)	
		PE-1OC12-SON-SMIR — (EOL)	
		PB-1OC12-SON-MM — (EOL)	M120, M320, T320, T640, T1600, TX Matrix
		PB-1OC12-SON-SMIR — (EOL)	
	4-port	PB-4OC12-SON-MM PB-4OC12-SON-SMIR	
OC12c/STM4, SFP	1-port	PB-1OC12-SON-SFP	M120, M320, T320, T640, T1600, TX Matrix
OC48c/STM16, SFP	1-port	PB-1OC48-SON-SFP	M120, M320, T320, T640, T1600, TX Matrix
		PB-1OC48-SON-B-SFP	
	4-port	PC-4OC48-SON-SFP	
OC192/STM64, XFP	1	PC-1OC192-SON-LR,	M320, T320, T640, T1600
		PC-1OC192-SON-SR2	
		PC-1OC192-VSR	
OC192/STM64, XFP	4	PD-4OC192-SON-XFP	M120, T640, T1600
OC768/STM256	1	PD-1OC768-SON-SR	T640, T1600

Fast Ethernet and Gigabit Ethernet PICs

Table 12 on page 211 lists the Fast Ethernet and Gigabit Ethernet PICs that are supported during a unified ISSU.



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

Table 12: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet

PIC Type	Number of Ports	Model Number	Router
Fast Ethernet	4	PB-4FE-TX	M120, M320, T320, T640, T1600, TX Matrix
		PE-4FE-TX	M10i
	8	PB-8FE-FX	M120, M320
		PE-8FE-FX	M10i
	12	PB-12FE-TX-MDI	M120, M320, T320
		PB-12FE-TX-MDIX	
		PE-12FE-TX-MDI	M10i
		PE-12FE-TX-MDIX	
	48	PB-48FE-TX-MDI	M120, M320, T320
		PB-48FE-TX-MDIX	
Gigabit Ethernet, SFP	1	PE-1GE-SFP	M10i
		PB-1GE-SFP	M120, M320, T320, T640, T1600, TX Matrix
	2	PB-2GE-SFP	
	4	PB-4GE-SFP	
	10	PC-10GE-SFP	
Gigabit Ethernet IQ, SFP	1	PE-1GE-SFP-QPP	M10i
		PB-1GE-SFP-QPP	M120, M320, T320, T640, T1600, TX Matrix
	2	PB-2GE-SFP-QPP	

Table 12: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet *(continued)*

PIC Type	Number of Ports	Model Number	Router
Gigabit Ethernet IQ2, SFP	4	PB-4GE-TYPE1-SFP-IQ2	M120, M320, T320, T640, T1600, TX Matrix
	8	PB-8GE-TYPE2-SFP-IQ2	
Gigabit Ethernet IQ2, XFP	8	PC-8GE-TYPE3-XFP-IQ2	M120, M320, T320, T640, T1600, TX Matrix
	1	PC-1XGE-TYPE3-XFP-IQ2	
10-Gigabit Ethernet, XENPAK	1	PC-1XGE-XENPAK	M120, M320, T320, T640, T1600, TX Matrix
10-Gigabit Ethernet, DWDM	1	PC-1XGE-DWDM-CBAND	M120, M320, T320, T640, T1600, TX Matrix
10-Gigabit Ethernet	4	PD-4XGE-XFP	T640, T1600, TX Matrix, TX Matrix Plus

Channelized PICs

Table 13 on page 212 lists the channelized PICs that are supported during a unified ISSU.

Table 13: Unified ISSU PIC Support: Channelized

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

Tunnel Services PICs

Table 14 on page 213 lists the tunnel services PICs that are supported during a unified ISSU.

Table 14: Unified ISSU PIC Support: Tunnel Services

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

ATM PICs

Table 15 on page 213 lists the ATM PICs that are supported during a unified ISSU. This includes support on Enhanced III FPCs.

Table 15: Unified ISSU PIC Support: ATM

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

Table 15: Unified ISSU PIC Support: ATM *(continued)*

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

Serial PICs

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

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

DS3, E1, E3, and T1 PICs

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

- 4-Port DS3 PIC (PB-4DS3)
- 4-Port E1 Coaxial PIC (PB-4E1-COAX)
- 4-Port E1 RJ48 PIC (PB-4E1-RJ48)
- 4-port E3 IQ PIC (PB-4E3-QPP)
- 4-Port T1 PIC (PB-4T1-RJ48)

Unified ISSU supports the following PICs on M10i routers:

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

Enhanced IQ PICs

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

- 1-port channelized OC12/STM4 enhanced IQ PIC (PB-1CHOC12-STM4-IQE-SFP)
- 1-port nonchannelized OC12/STM4 enhanced IQ PIC (PB-1OC12-STM4-IQE-SFP)
- 4-port channelized DS3/E3 enhanced IQ PIC (PB-4CHDS3-E3-IQE-BNC)
- 4-port nonchannelized DS3/E3 enhanced IQ PIC (PB-4DS3-E3-IQE-BNC)

Enhanced IQ2 Ethernet Services Engine (ESE) PIC

Unified ISSU supports the enhanced IQ2 ESE PICs listed in Table 16 on page 215.

Table 16: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC

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

Unified ISSU DPC and FPC Support on MX Series Routers

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

Chapter 21

Unified ISSU Configuration Guidelines

- Best Practices on page 217
- Before You Begin on page 217
- Performing a Unified ISSU on page 220
- Verifying a Unified ISSU on page 233
- Troubleshooting Unified ISSU Problems on page 233
- Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 234

Best Practices

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

- The master Routing Engine and backup Routing Engine must be running the same software version before you can perform a unified ISSU.
- During a unified ISSU, you cannot bring any PICs online or offline.

Before You Begin

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

1. Verify That the Master and Backup Routing Engines Are Running the Same Software Version on page 218
2. Back Up the Router Software on page 218
3. Verify That Graceful Routing Engine Switchover and Nonstop Active Routing Are Configured on page 219

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

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

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

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

Back Up the Router Software

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

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




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

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

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

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

```
{backup}

user@host> show system switchover

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

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

```
{master}

user@host> show task replication

Stateful Replication: Enabled
RE mode: Master

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

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

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

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

```
re0:
configuration check succeeds
re1:
commit complete
re0:
commit complete
```

Performing a Unified ISSU

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

1. Upgrading and Rebooting Both Routing Engines Automatically on page 220
2. Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually on page 225
3. Upgrading and Rebooting Only One Routing Engine on page 230

Upgrading and Rebooting Both Routing Engines Automatically

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

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

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

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

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

```
{backup}

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

```

re0:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite 9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.2]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]

```

```

re1:
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite [9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.20]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]

```

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

```

{master}

user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz reboot

ISSU: Validating Image
PIC 0/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no)
yes

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

```

```

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

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

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

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

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

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status

```

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

```

FPC 6           Online (ISSU)
FPC 7           Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/paKEuy' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by
PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

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

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

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

Saving package file in
/var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz ...
cp: /var/tmp/paKEuy is a directory (not copied).
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
Reboot consistency check bypassed - jinstall 9.0-20080114.2 will complete
installation upon reboot
[pid 30227]

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

System going down IMMEDIATELY

Connection to host closed.

```

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

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

```
{backup}

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

re0:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20080114.2]
JUNOS Base OS Software Suite 9.0-20080114.2]
JUNOS Kernel Software Suite [9.0-20080114.2]
JUNOS Crypto Software Suite [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20080114.2]
JUNOS Online Documentation [9.0-20080114.2]
JUNOS Routing Software Suite [9.0-20080114.2]

re1:
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0-20080114.2]
JUNOS Base OS Software Suite [9.0-20080114.2]
JUNOS Kernel Software Suite [9.0-20080114.2]
JUNOS Crypto Software Suite [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20080114.2]
JUNOS Online Documentation [9.0-20080114.2]
JUNOS Routing Software Suite [9.0-20080114.2]
```

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

```
{backup}

user@host> request chassis routing-engine master acquire

Attempt to become the master routing engine ? [yes,no] (no) yes

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

{master}
user@host>
```

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



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

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

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

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

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

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

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

```
{master}

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

re0:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
JUNOS Online Documentation [9.0R1]
JUNOS Routing Software Suite [9.0R1]

re1:
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
```

JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
 JUNOS Online Documentation [9.0R1]
 JUNOS Routing Software Suite [9.0R1]

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

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0R1.2-domestic-signed.tgz
```

```
ISSU: Validating Image
FPC 4 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no)
yes
```

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

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

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



```

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

```

```

Saving package file in /var/sw/pkg/jinstall-9.0R1.2-domestic-signed.tgz
...

```

```

Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

```

```

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover

```

```

Checking In-Service-Upgrade status

```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
FPC 4	Offline	Offlined by cli command
FPC 5	Online (ISSU)	

```

Resolving mastership...

```

```

Complete. The other routing engine becomes the master.

```

```

ISSU: RE switchover Done

```

```

ISSU: Upgrading Old Master RE

```

```

Installing package '/var/tmp/paeBi5' ...

```

```

Verified jinstall-9.0R1.2-domestic.tgz signed by PackageProduction_9_0_0

```

```

Adding jinstall...

```

```

Verified manifest signed by PackageProduction_9_0_0

```

```

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

```

```

Saving the config files ...

```

```

NOTICE: uncommitted changes have been saved in

```

```

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

```

```

Installing the bootstrap installer ...

```

```

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use
the
WARNING:      'request system reboot' command when software installation

```

```

is
WARNING:      complete. To abort the installation, do not reboot your
system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

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

```

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

```

{backup}

user@host> show version

re0:
-----
Hostname: user
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
JUNOS Online Documentation [9.0R1]
JUNOS Routing Software Suite [9.0R1]
Tabpkg [7.0]
JUNOS Installation Software [9.0R1.2]

re1:
-----
Hostname: user1
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
JUNOS Online Documentation [9.0R1.2]
JUNOS Routing Software Suite [9.0R1.2]

```

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

```
{backup}

user@host> request system reboot

Reboot the system ? [yes,no] (no) yes

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

{backup}
user@host>
System going down IMMEDIATELY

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

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

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

```
{backup}

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

re0:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
JUNOS Online Documentation [9.0R1.2]
JUNOS Routing Software Suite [9.0R1.2]

re1:
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
JUNOS Online Documentation [9.0R1.2]
JUNOS Routing Software Suite [9.0R1.2]
```

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

```

{backup}

user@host> request chassis routing-engine master acquire

Attempt to become the master routing engine ? [yes,no] (no) yes

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

{master}
user@host>

```

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



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

Upgrading and Rebooting Only One Routing Engine

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

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

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

```

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

```

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

```

{backup}

```

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

```
re0:
```

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

```
re1:
```

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

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

```
{master}
```

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz no-old-master-upgrade
```

```
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080116.2
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz
Verified jinstall-9.0-20080116.2-domestic.tgz signed by
PackageProduction_9_0_0
Using jinstall-9.0-20080116.2-domestic.tgz
Using jbundle-9.0-20080116.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpfe-9.0-20080116.2.tgz
Using jdocs-9.0-20080116.2.tgz
```

```

Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz'
...
Verified jinstall-9.0-20080116.2-domestic.tgz signed by
PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

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

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

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

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

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status

```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
FPC 5	Online (ISSU)	

```

Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE

{backup}
user@host>

```

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

Verifying a Unified ISSU

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

```

user@host> show chassis in-service-upgrade

```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
PIC 0	Online	
PIC 1	Online	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online	
PIC 1	Online	
FPC 5	Online	
PIC 0	Online	
FPC 6	Online	
PIC 3	Online	
FPC 7	Online	

For more information about the `show chassis in-service-upgrade` command, see the *JUNOS System Basics and Services Command Reference*.

Troubleshooting Unified ISSU Problems

If the unified ISSU procedure stops progressing, complete the following steps:

1. Open a new session on the master Routing Engine and issue the `request system software abort in-service-upgrade` command.
2. Check the existing router session to verify that the upgrade has been aborted.

An “ISSU: aborted!” message is provided. Additional system messages provide you with information about where the upgrade stopped and recommendations for the next step to take.

For more information about the `request system software abort in-service-upgrade` command, see the *JUNOS System Basics and Services Command Reference*.

Managing and Tracing BFD Sessions During Unified ISSU Procedures

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

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

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

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



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

For more information about BFD, see the *JUNOS Routing Protocols Configuration Guide*.

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

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


Chapter 22

Unified ISSU Configuration Statements Summary

This chapter provides a reference for each of the unified in-service software upgrade (ISSU) configuration statements. The statements are organized alphabetically.



NOTE: To perform a unified ISSU, you must first configure graceful Routing Engine switchover and nonstop active routing (NSR).

no-issu-timer-negotiation

Syntax no-issu-timer-negotiation;

Hierarchy Level [edit protocols bfd],
[edit logical-systems *logical-system-name* protocols bfd],
[edit routing-instances *routing-instance-name* protocols bfd]

Release Information Statement introduced in JUNOS Release 9.1.

Description Disable unified ISSU timer negotiation for Bidirectional Forwarding Detection (BFD) sessions.



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

Usage Guidelines See “Managing and Tracing BFD Sessions During Unified ISSU Procedures” on page 234.

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

Related Topics For details on BFD, see the *JUNOS Routing Protocols Configuration Guide*.

traceoptions

Syntax traceoptions {
 file *name* <size *size*> <files *number*> <world-readable | no-world-readable>;
 flag *flag* <flag-modifier> <disable>;
 }

Hierarchy Level [edit protocols bfd]

Release Information Statement introduced before JUNOS Release 7.4.
 issu flag for BFD added in JUNOS Release 9.1.

Description Define tracing operations that track unified in-service software upgrade (ISSU) functionality in the router.

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

Default If you do not include this statement, no global tracing operations are performed.

Options disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *name*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

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.

Range: 2 through 1000 files

Default: 2 files

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

flag *flag*—Tracing operation to perform. There is only one unified ISSU tracing option:

- **issu**—Trace BFD unified ISSU operations.

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.

Usage Guidelines See “Managing and Tracing BFD Sessions During Unified ISSU Procedures” on page 234.

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

Part 9

Index

- Index on page 241
- Index of Statements and Commands on page 247

Index

Symbols

#, comments in configuration statements.....	xxviii
(), in syntax descriptions.....	xxviii
< >, in syntax descriptions.....	xxviii
[], in configuration statements.....	xxviii
{ }, in configuration statements.....	xxviii
(pipe), in syntax descriptions.....	xxviii

A

accept-data statement.....	175
usage guidelines.....	164
advertise-interval statement.....	176
usage guidelines.....	162
advertisement intervals, VRRP.....	162
aggregate routes	
graceful restart.....	95
authentication, VRRP.....	161
authentication-key statement.....	177
usage guidelines.....	161
authentication-type statement.....	178
usage guidelines.....	161

B

backup routers, VRRP.....	155
bandwidth-threshold statement.....	179
usage guidelines.....	165
BFD, nonstop active routing.....	76
BGP	
graceful restart.....	95
nonstop active routing.....	74
unified ISSU.....	206
bidirectional forwarding detection <i>See</i> BFD	
Border Gateway Protocol <i>See</i> BGP	
braces, in configuration statements.....	xxviii
brackets	
angle, in syntax descriptions.....	xxviii
square, in configuration statements.....	xxviii

C

CCC, graceful restart.....	95
----------------------------	----

CFEB redundancy	
configuring.....	28
overview.....	16
cfeb statement.....	33
usage guidelines.....	28
circuit cross-connect <i>See</i> CCC	
comments, in configuration statements.....	xxviii
commit synchronize statement.....	88
usage guidelines.....	82
Compact Forwarding Engine Board <i>See</i> CFEB	
redundancy	
configuration files, copying between Routing Engines.....	23
conventions	
text and syntax.....	xxvii
curly braces, in configuration statements.....	xxviii
customer support.....	xxix
contacting JTAC.....	xxix

D

description statement.....	34
usage guidelines.....	29
disable statement.....	141
usage guidelines	
aggregate routes.....	103
BGP.....	104
ES-IS.....	105
global.....	104, 109, 111, 112
IS-IS.....	105
OSPF	106
OSPFv3.....	106
PIM.....	107
RIP.....	107
RIPng.....	107
routing instances.....	111
routing instances inside a logical system.....	112
RSVP.....	109
static routes.....	103
distributed ppmv process.....	170
documentation	
comments on.....	xxviii

E

ES-IS	
graceful restart	95

F

failover on-disk-failure statement	34
usage guidelines	26
failover on-loss-of-keepalives statement	35
usage guidelines	26
failover other-routing-engine statement	36
usage guidelines	27
fast-interval statement	180
usage guidelines	163
FEB redundancy	
configuration example	30
configuring	29
overview	16
feb statement	
edit chassis redundancy hierarchy	37
redundancy group	37
usage guidelines	29
file copy command	
usage guidelines	23
font conventions	xxvii
Forwarding Engine Board	<i>See</i> FEB redundancy

G

graceful restart	
commands, operational mode	113
concepts	95
configuration procedure	
aggregate routes	103
MPLS protocols	108
routing protocols	103
static routes	103
VPNs	111
overview	95
aggregate routes	96
MPLS protocols	98
routing protocols	96
static routes	96
VPNs	100
protocol support	95
sample configuration	115
system requirements	96
trace options	108
verifying operation of	113
graceful Routing Engine switchover	
concepts	45
DPC support	50
enabling	53
feature support	49
PIC support	51
platform support	49

subscriber access support	51
system architecture	46
system requirements	48
verifying status of	54
graceful-restart statement	142
usage guidelines	
aggregate routes	103
BGP	104
global	104, 109, 111, 112
IS-IS	105
LDP	110
OSPF	106
OSPFv3	106
PIM	107
RIP	107
RIPng	107
routing instances	111
routing instances inside a logical system	112
RSVP	109
static routes	103
graceful-switchover statement	57
usage guidelines	53
GRES	<i>See</i> graceful Routing Engine switchover

H

helper-disable statement	143
usage guidelines	
IS-IS	105
LDP	110
OSPF	106
OSPFv3	106
RSVP	109
high availability features	
graceful restart	95
graceful Routing Engine switchover	45
nonstop active routing	71
nonstop bridging	61
overview	3
unified ISSU	199
VRRP	155
hold-time statement	181

I

icons defined, notice	xxvii
in-service software upgrade	<i>See</i> unified ISSU
inet6-advertise-interval statement	182
usage guidelines	163
initial configuration, redundant Routing Engines	22
interface statement	183
usage guidelines	165
Intermediate System-to-Intermediate System	<i>See</i> IS-IS

IS-IS	
graceful restart.....	95
nonstop active routing.....	74
unified ISSU.....	206
ISSU <i>See</i> unified ISSU	

K

keepalive-time statement.....	38
usage guidelines.....	26

L

L2CKT, nonstop active routing.....	77
Label Distribution Protocol <i>See</i> LDP	
LACP	
unified ISSU.....	207
Layer 2 circuits	
nonstop active routing.....	74, 77
unified ISSU.....	206
Layer 2 circuits, nonstop active routing.....	74, 77
Layer 2 VPNs	
nonstop active routing.....	74
Layer 3 VPNs	
nonstop active routing.....	74
unified ISSU.....	206
LDP	
graceful restart.....	95
nonstop active routing.....	74
unified ISSU.....	206
LDP-based VPLS, nonstop active routing.....	77

M

manuals	
comments on.....	xxviii
master router, VRRP.....	155
maximum-helper-recovery-time statement.....	143
usage guidelines.....	109
maximum-helper-restart-time statement.....	144
usage guidelines.....	109
maximum-neighbor-reconnect-time	
usage guidelines.....	110
maximum-neighbor-reconnect-time statement.....	144
maximum-neighbor-recovery-time statement.....	145
usage guidelines.....	110
MPLS, graceful restart.....	95
MSTP, nonstop bridging.....	64
Multiple Spanning Tree Protocol <i>See</i> MSTP	

N

next-generation RIP <i>See</i> RIPng	
no-accept-data statement.....	183
usage guidelines.....	164

no-auto-failover statement.....	39
usage guidelines.....	29
no-issu-timer-negotiation statement.....	235
usage guidelines.....	234
no-preempt statement.....	183
usage guidelines.....	163
no-strict-lsa-checking statement.....	145
usage guidelines	
OSPF.....	106
OSPFv3.....	106
nonstop active routing	
concepts.....	71
enabling.....	81
platform support.....	74
protocol support.....	74
sample configuration.....	84
system architecture.....	72
system requirements.....	73
trace options.....	83
verifying status of	82
nonstop bridging	
concepts.....	61
enabling.....	65
platform support.....	63
protocol support.....	64
system architecture.....	62
system requirements.....	63
verifying status of	66
nonstop-bridging statement.....	67
usage guidelines.....	65
nonstop-routing statement.....	89
usage guidelines.....	81
notice icons defined.....	xxvii
notify-duration statement.....	146
usage guidelines.....	106
NSR <i>See</i> nonstop active routing	

O

Open Shortest Path First <i>See</i> OSPF	
OSPF	
graceful restart.....	95
nonstop active routing.....	74
unified ISSU.....	206
OSPFv3	
nonstop active routing.....	74
unified ISSU.....	206

P

parentheses, in syntax descriptions.....	xxviii
passive ARP learning, VRRP.....	169
periodic packet management process.....	170
PIM	
unified ISSU.....	206
PIM, graceful restart.....	95

PIM, nonstop active routing.....	77
ppmd process.....	170
preempt statement.....	184
usage guidelines.....	163
preempting master router, VRRP.....	163
priority statement.....	185
usage guidelines.....	159
priority-cost statement.....	186
usage guidelines.....	165
priority-hold-time statement.....	187
usage guidelines.....	165
Protocol Independent Multicast <i>See</i> PIM	

R

Rapid Spanning Tree Protocol <i>See</i> RSTP	
reconnect-time statement.....	146
usage guidelines.....	110
recovery-time statement.....	147
usage guidelines.....	110
redundancy feb statement	
usage guidelines.....	29
redundancy statement.....	39
usage guidelines.....	21
redundancy-group statement.....	40
usage guidelines.....	29
request chassis cfeb master switch command	
usage guidelines.....	28
request chassis redundancy feb command	
usage guidelines.....	29
request chassis routing-engine master acquire command	
usage guidelines.....	27
request chassis routing-engine master release command	
usage guidelines.....	27
request chassis routing-engine master switch command	
usage guidelines.....	27
request chassis sfm master switch command	
usage guidelines.....	31
request chassis ssb master switch command	
usage guidelines.....	31
request routing-engine login command	
usage guidelines.....	23
request system software add command	
usage guidelines.....	24
Resource Reservation Protocol <i>See</i> RSVP	
restart-duration statement.....	148
usage guidelines	
ES-IS.....	105
global.....	104, 109, 111, 112
IS-IS.....	105
OSPF	106
OSPFv3.....	106
PIM.....	107

routing instances.....	111
routing instances inside a logical system.....	112
restart-time statement.....	149
usage guidelines	
BGP.....	104
RIP.....	107
RIPng.....	107
RIP	
graceful restart.....	95
nonstop active routing.....	74
RIPng	
graceful restart.....	95
nonstop active routing.....	74
route statement.....	188
Routing Engine redundancy	
copying configuration files.....	23
default behavior.....	13
failover	
conditions.....	12
on loss of keepalive signal.....	26
graceful Routing Engine switchover.....	25
halting Routing Engines.....	15
initial configuration.....	22
log file, viewing.....	27
mastership	
default setup, modifying.....	25
switching, manually.....	27
overview.....	11
software packages, loading	24
TX Matrix, running the same JUNOS release.....	14
Routing Engine switchover effects	
comparison of high availability features.....	6, 48
Routing Information Protocol <i>See</i> RIP	
Routing Information Protocol next generation <i>See</i> RIPng	
routing-engine statement.....	40
usage guidelines.....	25
RSTP, nonstop bridging.....	64
RSVP	
graceful restart.....	95

S

set delegate-processing statement.....	170
SFM redundancy	
configuring.....	31
overview.....	19
sfm statement.....	41
usage guidelines.....	31
show bgp replication command	
usage guidelines.....	82
show chassis cfeb command	
usage guidelines.....	28
show chassis feb command	
usage guidelines.....	29

show chassis sfm command
 usage guidelines.....31

show chassis ssb command
 usage guidelines.....32

show task replication command
 usage guidelines.....82

software packages
 transferring between Routing Engines.....24

Spanning Tree Protocol *See* STP

SSB redundancy
 configuring.....31
 overview.....18

ssb statement.....41
 usage guidelines.....31

stale-routes-time statement.....150
 usage guidelines.....104

startup period, VRRP.....159

startup-silent-period statement.....188
 usage guidelines.....159

static routes
 graceful restart.....95

STP, nonstop bridging.....64

support, technical *See* technical support

Switching and Forwarding Module *See* SFM redundancy

switching control board redundancy *See* CFEB
 redundancy, FEB redundancy, SFM redundancy, SSB
 redundancy

synchronizing Routing Engines
 graceful Routing Engine switchover.....54
 nonstop active routing.....82
 Routing Engine redundancy.....65

syntax conventions.....xxvii

System and Switch Board *See* SSB redundancy

system requirements
 graceful restart.....96
 graceful Routing Engine switchover.....48
 nonstop active routing.....73
 nonstop bridging.....63
 unified ISSU.....205

T

TCC, graceful restart.....95

technical support
 contacting JTAC.....xxix

traceoptions statement
 graceful restart151
 usage guidelines.....108

nonstop active routing.....90
 usage guidelines.....83

unified ISSU.....236

VRRP.....189
 usage guidelines.....168

track statement.....191
 usage guidelines.....165

translational cross-connect *See* TCC

TX Matrix router
 Routing Engine redundancy.....14

U

unified in-service software upgrade *See* unified ISSU

unified ISSU
 concepts.....199
 configuration procedure
 best practices.....217
 performing an.....220
 preparing for.....217

DPC support.....215

PIC support
 ATM.....213
 channelized.....212
 DS3.....214
 E1.....214
 E3 IQ.....214
 Fast Ethernet.....211
 FPC/PIC compatibility.....208
 Gigabit Ethernet.....211
 restrictions.....209
 serial.....214
 SONET/SDH.....209
 T1.....214
 tunnel services.....213

platform support.....206

protocol support.....206

system requirements.....205

troubleshooting.....233

verifying status of.....233

V

Virtual Router Redundancy Protocol *See* VRRP

virtual-address statement.....192
 usage guidelines.....159

virtual-inet6-address statement.....192
 usage guidelines.....159

virtual-link-local-address statement.....193
 usage guidelines.....159

VPLS
 nonstop active routing.....74

VPNs, graceful restart.....95

VRRP
 advertisement interval.....162
 authentication.....161
 basic configuration.....159
 group
 inheritance.....167

overview.....155

passive ARP learning.....169

preempting master router.....163

sample configuration.....171

trace operations.....	168
tracking logical interface status.....	165
vrrp-group statement.....	194
usage guidelines.....	159
vrrp-inet6-group statement.....	195
usage guidelines.....	159
vrrpd process.....	170

Index of Statements and Commands

A

accept-data statement.....	175
advertise-interval statement.....	176
authentication-key statement.....	177
authentication-type statement.....	178

B

bandwidth-threshold statement.....	179
------------------------------------	-----

C

cfeb statement.....	33
commit synchronize statement.....	88

D

description statement.....	34
disable statement.....	141

F

failover on-disk-failure statement.....	34
failover on-loss-of-keepalives statement	35
failover other-routing-engine statement.....	36
fast-interval statement.....	180
feb statement	
edit chassis redundancy hierarchy.....	37
redundancy group.....	37

G

graceful-restart statement.....	142
graceful-switchover statement.....	57

H

helper-disable statement.....	143
hold-time statement.....	181

I

inet6-advertise-interval statement.....	182
interface statement.....	183

K

keepalive-time statement.....	38
-------------------------------	----

M

maximum-helper-recovery-time statement.....	143
maximum-helper-restart-time statement.....	144
maximum-neighbor-reconnect-time statement.....	144
maximum-neighbor-recovery-time statement.....	145

N

no-accept-data statement.....	183
no-auto-failover statement.....	39
no-issu-timer-negotiation statement.....	235
no-preempt statement.....	183
no-strict-lsa-checking statement.....	145
nonstop-bridging statement.....	67
nonstop-routing statement.....	89
notify-duration statement.....	146

P

preempt statement.....	184
priority statement.....	185
priority-cost statement.....	186
priority-hold-time statement.....	187

R

reconnect-time statement.....	146
recovery-time statement.....	147
redundancy statement.....	39
redundancy-group statement.....	40
restart-duration statement.....	148
restart-time statement.....	149
route statement.....	188
routing-engine statement.....	40

S

set delegate-processing statement.....	170
sfm statement.....	41
ssb statement.....	41

stale-routes-time statement.....	150
startup-silent-period statement.....	188

T

traceoptions statement	
graceful restart.....	151
nonstop active routing.....	90
unified ISSU.....	236
VRRP.....	189
track statement.....	191

V

virtual-address statement.....	192
virtual-inet6-address statement.....	192
virtual-link-local-address statement.....	193
vrrp-group statement.....	194
vrrp-inet6-group statement.....	195