



JUNOS[™]e Software for E Series[™] Broadband Services Routers

Service Availability Configuration Guide

Release 10.3.x

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Published: 2009-10-07

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.

JUNOSe™ Software for E Series™ Broadband Services Routers Service Availability Configuration Guide

Release 10.3.x

Copyright © 2009, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing: Krupa Chandrashekar, Sairam Venugopalan

Editing: Benjamin Mann

Illustration: Nathaniel Woodward

Cover Design: Edmonds Design

Revision History

October 2009— FRS JUNOSe 10.3.x

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 NON-INFRINGEMENT. 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 the Documentation	xix
Part 1	Chapters	
Chapter 1	Service Availability	3
Chapter 2	Managing Module Redundancy	7
Chapter 3	Managing Stateful SRP Switchover	25
Chapter 4	Configuring a Unified In-Service Software Upgrade	55
Chapter 5	Configuring VRRP	101
Chapter 6	Managing Interchassis Redundancy	121
Part 2	Index	
	Index	143

Table of Contents

	About the Documentation	xix
	E Series and JUNOS ^e Documentation and Release Notes	xix
	Audience	xix
	E Series and JUNOS ^e Text and Syntax Conventions	xix
	Obtaining Documentation	xxi
	Documentation Feedback	xxi
	Requesting Technical Support	xxi
	Self-Help Online Tools and Resources	xxii
	Opening a Case with JTAC	xxii
Part 1	Chapters	
Chapter 1	Service Availability	3
	Service Availability Overview	3
	Service Availability Versus High Availability	4
	Understanding Service Availability Features	5
	5
	Module Redundancy	5
	Stateful SRP Switchover	5
	Unified ISSU	5
	VRRP	6
	Interchassis Redundancy	6
Chapter 2	Managing Module Redundancy	7
	Line Module Redundancy Overview	7
	Module Requirements	7
	ERX7xx Models and ERX14xx Models	7
	E120 and E320 Routers	8
	Automatic Switchover	9
	Limitations of Automatic Switchover	9
	Reversion After Switchover	9
	Configuring Line Module Redundancy	10
	Managing Line Module Redundancy	10
	SRP Module Redundancy	11
	SRP Module Behavior	11
	Specifying the Configuration for Redundant SRP Modules	14

Installing a Redundant SRP Module	15
Managing SRP Module Redundancy	16
Switching to the Redundant SRP Module	18
Upgrading Software on a Redundant SRP Module	19
Monitoring the Status LEDs	19
Monitoring Line Module and SRP Module Redundancy	19
Managing Port Redundancy	23

Chapter 3**Managing Stateful SRP Switchover 25**

Understanding Stateful SRP Switchover	25
Platform Considerations	26
Module Requirements	26
Redundancy Modes of Operation	26
File System Synchronization Mode	26
High Availability Mode	27
Understanding SRP State Behavior	28
Disabled State	29
Initializing State	29
Active State	30
Pending State	31
Application Support	31
Before Activating High Availability	40
Activating High Availability	40
Deactivating High Availability	41
Setting the IP Interface Priority	42
Upgrading Software	42
Monitoring Stateful SRP Switchover	43
Stateful SRP Switchover show Commands	43
Clearing the Redundancy History	52

Chapter 4**Configuring a Unified In-Service Software Upgrade 55**

Unified ISSU Overview	55
Router Behavior During a Unified In-Service Software Upgrade	56
Unified ISSU Platform Considerations	57
Unified ISSU Terms That Describe SRP and Line Module Behavior	57
Unified ISSU References	58
Unified ISSU Phases Overview	58
Unified ISSU Initialization Phase Overview	59
Application Data Upgrade on the Standby SRP Module	60
SNMP Traps	60
Unified ISSU Upgrade Phase Overview	60
Exceptions During the Upgrade Phase	62
Verification of Requirements	62
Upgrade Setup	63
Unified ISSU Service Restoration Phase Overview	65
Application Support for Unified ISSU	65

Unexpected Application-Specific Behavior During Unified ISSU	73
AAA Authentication and Authorization Disabled	74
ATM Affected Behaviors	74
ILMI Sessions Not Maintained	74
OAM CC Effects on VCC	74
OAM VC Integrity Verification Cessation	74
Port Data Rate Monitoring Cessation	75
VC and VP Statistics Monitoring Halts Unified ISSU Progress	75
DHCP Affected Behaviors	75
DHCP Common Component Information Suspended	75
DHCP Relay and DHCP Relay Proxy Prevent Unified ISSU	75
DHCP Packet Capture Halted on Line Modules	75
DoS Protection State Freeze	75
Ethernet Affected Behaviors	76
ARP Packets Briefly Not Sent or Received	76
Link Aggregation interruption	76
Port Data Rate Monitoring Halted	76
VLAN Statistics Monitoring Halts Unified ISSU Progress	77
FTP Server File Transfer Behaviors	77
IS-IS Effects on Graceful Restart and Network Stability	79
Configuring Graceful Restart Before Unified ISSU Begins	80
Configuring Graceful Restart When BGP And LDP Are Configured	80
Routing Around the Restarting Router to Minimize Network Instability	80
L2TP Failover of Established Tunnels	81
OSPF Effects on Graceful Restart, Timeouts, and Network Stability	82
Configuring Graceful Restart Before Unified ISSU Begins	82
Configuring Graceful Restart When BGP And LDP Are Configured	82
Configuring a Longer Dead Interval Than Normal	83
Routing Around the Restarting Router to Minimize Network Instability	83
PIM Suspended During Unified ISSU	84
Subscriber Logins and Logouts Suspended During Unified ISSU	84
Subscriber Statistics Accumulation or Deletion	84
SONET/SDH Behavior During Unified ISSU	85
T3	85
TACACS+ Services Not Available	85
Interruption in Traffic Forwarding for Layer 3 Routing and Signaling Protocols	86
Recommended Routing Protocol Timer Settings	88
Before You Begin a Unified In-Service Software Upgrade	89
Hardware Requirements for Unified ISSU	89
Software Requirements for Unified ISSU	90
Upgrading Router Software with Unified ISSU	91
Halting the Unified ISSU Process and Restoring the Original State of the Router	95
Halting Unified ISSU During Initialization Phase	95
Halting Unified ISSU During Upgrade Phase	96
Monitoring a Unified In-Service Software Upgrade	97

Chapter 5	Configuring VRRP	101
	VRRP Overview	101
	VRRP Terms	101
	Platform Considerations	102
	References	102
	How VRRP Works	103
	Configuration Examples	103
	Basic VRRP Configuration	103
	Commonly Used VRRP Configuration	104
	VRRP Configuration Without the Real Address Owner	105
	How VRRP Is Implemented in E Series Routers	106
	Router Election Rules	107
	Configuring VRRP	108
	Configuring the IP Interface	108
	Creating VRIDs	108
	Configuration Steps	109
	Changing Object Priority	112
	Monitoring VRRP	113
 Chapter 6	 Managing Interchassis Redundancy	 121
	ICR Overview	121
	ICR Platform Considerations	123
	Interface Specifiers	123
	ICR Terms	124
	ICR References	124
	ICR Scaling Considerations	124
	1:1 Subscriber Redundancy in a 4-Node ICR Cluster	125
	1:3 Subscriber Redundancy in a 4-Node ICR Cluster	125
	Guidelines for Deploying an ICR Partition in Your Network	126
	Hardware Requirements for ICR	127
	Network Requirements for ICR	127
	Router Configurations for ICR	127
	Interaction with RADIUS for ICR	127
	ICR Partition Accounting Overview	128
	Configuring an ICR Partition	129
	Configuring the Interface on Which the ICR Partition Resides	130
	Configuring VRRP Instances to Match ICR Requirements	131
	Naming ICR Partitions	131
	Grouping ICR Subscribers Based on S-VLAN IDs	132
	Grouping ICR Subscribers Based on VLAN IDs	133
	Example: Configuring ICR Partitions That Group Subscribers by S-VLAN ID	134
	Using RADIUS to Manage Subscribers Logging In to ICR Partitions	136
	Monitoring the Configuration of an ICR Partition Attached to an Interface	137
	Monitoring the Configuration of ICR Partitions	138

Part 2**Index**

Index	143
-------------	-----

List of Figures

Part 1

Chapters

Chapter 2	Managing Module Redundancy	7
	Figure 1: SRP Module on ERX7xx Models and ERX14xx Models	13
	Figure 2: SRP Module on the E120 and E320 Routers	14
Chapter 3	Managing Stateful SRP Switchover	25
	Figure 3: High Availability States	28
Chapter 5	Configuring VRRP	101
	Figure 4: Basic VRRP Configuration	104
	Figure 5: Commonly Used VRRP Configuration	105
	Figure 6: VRRP Configuration Without the Real Address Owner	106
Chapter 6	Managing Interchassis Redundancy	121
	Figure 7: Sample Network for ICR Deployment	122
	Figure 8: Sample 1:1 Subscriber Redundancy in a 4-Node ICR Cluster	125
	Figure 9: Sample 1:3 Subscriber Redundancy in a 4-Node ICR Cluster	126

List of Tables

About the Documentation	xix
Table 1: Notice Icons	xx
Table 2: Text and Syntax Conventions	xx

Part 1

Chapters

Chapter 2	Managing Module Redundancy	7
	Table 3: Commands That Can Cause Automatic Switchover	9
	Table 4: Function of the Online and Redundant LEDs	19
Chapter 3	Managing Stateful SRP Switchover	25
	Table 5: Application Support for Stateful SRP Switchover	32
Chapter 4	Configuring a Unified In-Service Software Upgrade	55
	Table 6: Unified ISSU-Related Terms	57
	Table 7: Router Response to Undesirable Events During the Upgrade Phase	62
	Table 8: Application Support for Unified In-Service Software Upgrades	66
	Table 9: Behavior of Routing Protocols During a Unified In-Service Software Upgrade	87
	Table 10: Recommended Routing Protocol Timer Settings	88
Chapter 5	Configuring VRRP	101
	Table 11: VRRP Definitions	102
Chapter 6	Managing Interchassis Redundancy	121
	Table 12: ICR Terminology	124
	Table 13: show icr-partition Output Fields	137
	Table 14: show icr-partitions Output Fields	139

About the Documentation

- E Series and JUNOSe Documentation and Release Notes on page xix
- Audience on page xix
- E Series and JUNOSe Text and Syntax Conventions on page xix
- Obtaining Documentation on page xxi
- Documentation Feedback on page xxi
- Requesting Technical Support on page xxi

E Series and JUNOSe Documentation and Release Notes

For a list of related JUNOSe documentation, see
<http://www.juniper.net/techpubs/software/index.html>.

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/>.

Audience

This guide is intended for experienced system and network specialists working with Juniper Networks E Series Broadband Services Routers in an Internet access environment.

E Series and JUNOSe Text and Syntax Conventions

Table 1 on page xx defines notice icons used in this documentation.

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 xx defines text and syntax conventions that we use throughout the E Series and JUNOS documentation.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents commands and keywords in text.	<ul style="list-style-type: none"> Issue the clock source command. Specify the keyword exp-msg.
Bold text like this	Represents text that the user must type.	host1(config)# traffic class low-loss1
Fixed-width text like this	Represents information as displayed on your terminal's screen.	<pre>host1#show ip ospf 2 Routing Process OSPF 2 with Router ID 5.5.0.250 Router is an Area Border Router (ABR)</pre>
<i>Italic text like this</i>	<ul style="list-style-type: none"> Emphasizes words. Identifies variables. Identifies chapter, appendix, and book names. 	<ul style="list-style-type: none"> There are two levels of access: <i>user</i> and <i>privileged</i>. <i>clusterId</i>, <i>ipAddress</i>. <i>Appendix A, System Specifications</i>
Plus sign (+) linking key names	Indicates that you must press two or more keys simultaneously.	Press Ctrl + b.
Syntax Conventions in the Command Reference Guide		
Plain text like this	Represents keywords.	terminal length
<i>Italic text like this</i>	Represents variables.	<i>mask</i> , <i>accessListName</i>
(pipe symbol)	Represents a choice to select one keyword or variable to the left or to the right of this symbol. (The keyword or variable can be either optional or required.)	diagnostic line

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

Convention	Description	Examples
[] (brackets)	Represent optional keywords or variables.	[internal external]
[]* (brackets and asterisk)	Represent optional keywords or variables that can be entered more than once.	[level1 level2 l1]*
{ } (braces)	Represent required keywords or variables.	{ permit deny } { in out } { <i>clusterId</i> <i>ipAddress</i> }

Obtaining Documentation

To obtain the most current version of all Juniper Networks technical documentation, see the products documentation page on the Juniper Networks Web site at <http://www.juniper.net/>.

To download complete sets of technical documentation to create your own documentation CD-ROMs or DVD-ROMs, see the CD-ROM and DVD-ROM Documentation page at

<http://www.juniper.net/techpubs/resources/cdrom.html>

Copies of the Management Information Bases (MIBs) available in a software release are included on the software CDs and at <http://www.juniper.net/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation to better meet your needs. 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

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <http://www.juniper.net/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, see http://www.juniper.net/support/requesting_support.html .

Part 1

Chapters

- Service Availability on page 3
- Managing Module Redundancy on page 7
- Managing Stateful SRP Switchover on page 25
- Configuring a Unified In-Service Software Upgrade on page 55
- Configuring VRRP on page 101
- Managing Interchassis Redundancy on page 121

Chapter 1

Service Availability

This chapter explains what is service availability and discusses the features of service availability. It also discusses Juniper Networks multi-layered service availability approach for uninterrupted delivery of services.

- Service Availability Overview on page 3
- Understanding Service Availability Features on page 5

Service Availability Overview

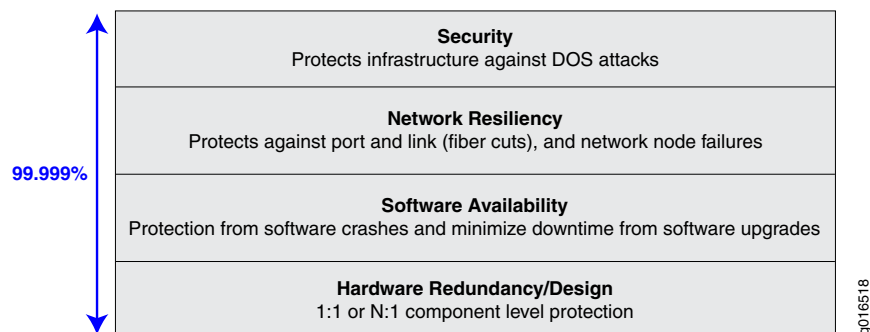
In a conventional network, router outages can occur because of Denial of Service (DoS) attacks, line module or SRP module failure, software defects, feature upgrades, or whole router failure. These outages result in subscriber downtime.

To reduce subscriber downtime, a network must have the following capabilities:

- *Reliability* — A network that does not crash often and recovers from failure very rapidly. During recovery, the network maintains user sessions and forwards data with little or no impact on the delivery of services.
- *Resiliency* — A network component or network's that responds to failure, resists failure, and handles failure with little or no impact on the delivery of services.
- *Redundancy* — A network whose reliability is enhanced by the addition of a backup component.
- *High Availability* — A network that is both reliable and resilient.

JUNOS software uses a multi-layered service availability approach that enables you to provide uninterrupted delivery of services with the help of reliable, highly available, and redundant hardware and software components.

Figure 1 illustrates the multiple layers of JUNOS Software service availability



The security layer protects the network from DoS attacks.

The network resiliency layer protects against port, link, and node failures. You can configure IEEE 802.3ad link aggregation for Ethernet, and Virtual Router Redundancy Protocol (VRRP) to improve network resiliency.

The software availability layer protects against software failures by using hot-fixes or installing a higher numbered software release. You can perform a unified in-service software upgrade (ISSU) instead of the conventional software upgrade to reduce outage. You can eliminate or reduce single points of failure by configuring stateful SRP switchover (high availability). Any network component with an uptime of 99.999 percent is considered *highly available* with a downtime of less than five minutes in a year.

The hardware redundancy and design layer introduces redundancy in the network in the form of multiple power supplies, cooling devices, line modules, and sometimes even a router. For instance, you can install a backup line module in your router to protect against line module failure. You can also configure a router as a backup router that accepts subscriber login requests when the master router fails.

Service Availability Versus High Availability

High availability is a measure of the uptime of a network or network component. A network component that has a downtime of 5 minutes is accessible or available 99 percent of the time. If a failure occurs, a backup component is available within 5 minutes. A highly available network is a network that has components that either have high reliability or have the ability to recover very quickly from a failure, or both.

Service availability refers to the ability to provide uninterrupted delivery of services. For example, from the time when the component failed to the time when the backup component was accessible, the delivery of services is interrupted. To provide uninterrupted delivery of services, highly available components must maintain session details and other data across failures. Service availability can thus be defined as the ability to provide uninterrupted delivery of services using a highly available network.

Related Topics ■ Understanding Service Availability Features on page 5

Understanding Service Availability Features

Service availability refers to ability of a network or a network component to provide uninterrupted delivery of services using highly available, redundant, and reliable components. This topic provides brief overviews of the benefits of using the following service availability features:

- Module Redundancy on page 5
- Stateful SRP Switchover on page 5
- Unified ISSU on page 5
- VRRP on page 6
- Interchassis Redundancy on page 6

Module Redundancy

For hardware components, Juniper Networks provides redundancy solutions to ensure that the router continues to operate in the event of a hardware fault. Redundancy also enables you to hot-swap various components within your E Series router.

Stateful SRP Switchover

Stateful SRP switchover (high availability) enables you to reduce or eliminate single points of failure in your network. Stateful SRP switchover provides both hardware-specific and software-specific methods to ensure minimal downtime and ultimately improve the performance of your network.

Stateful SRP switchover minimizes the impact to the router of a stateful switchover from the active SRP module to the standby SRP module. Stateful SRP switchover maintains user sessions and data forwarding through the router during the switchover, thus improving the overall availability of the router.

Unified ISSU

A conventional software upgrade—one that does not use the unified in-service software upgrade (ISSU) process—causes a router-wide outage for all users. Only static configurations (stored on the flash card) are maintained across the upgrade; all dynamic configurations are lost. A conventional upgrade can take 30-40 minutes to complete, with additional time required to bring all users back online.

Unified ISSU enables you to upgrade the router to a higher-numbered software release without disconnecting user sessions or disrupting forwarding through the chassis.

When an application supports unified ISSU, you can configure the application on the router and proceed with the unified in-service software upgrade with no adverse effect on the upgrade.

When you perform a unified ISSU on a router that has one or more modules that do not support unified ISSU, these modules are upgraded by means of the legacy, conventional upgrade process. The unsupported modules undergo a cold reboot at the beginning of the unified ISSU process, and are held down until the ISSU process is completed.

VRRP

Virtual Router Redundancy Protocol (VRRP) prevents loss of network connectivity to end hosts when the static default IP gateway fails. By implementing VRRP, you can designate a number of routers as backup routers in the event that the default master router fails. In case of a failure, VRRP dynamically shifts the packet-forwarding responsibility to a backup router. VRRP creates a redundancy scheme that enables hosts to keep a single IP address for the default gateway but maps the IP address to a well-known virtual MAC address. You can take advantage of the redundancy provided by VRRP without performing any special configuration on the end host systems.

Routers running VRRP dynamically elect master and backup routers. You can also force assignment of master and backup routers using priorities in the range 1 – 255, with 255 being the highest priority.

VRRP supports virtual local area networks (VLANs), stacked VLANs (S-VLANs), and creation of interchassis redundancy (ICR) partitions.

Interchassis Redundancy

ICR enables you to minimize subscriber downtime when the chassis or access interface on the edge router fails by re-creating subscriber sessions that were originally terminated on the failed router. It also enables you to track the failure of uplink interfaces. In this way, ICR enables you to completely recover from router failure. By using extended VRRP features, ICR enables you to track the failure of uplink interfaces, elect the master VRRP instance, and detect failure of a VRRP instance.

- Related Topics**
- *Chapter 2, Managing Module Redundancy*
 - *Chapter 3, Managing Stateful SRP Switchover*
 - *Chapter 4, Configuring a Unified In-Service Software Upgrade*
 - *Chapter 5, Configuring VRRP*
 - *Chapter 6, Managing Interchassis Redundancy*
 - *Service Availability Overview on page 3*

Chapter 2

Managing Module Redundancy

This chapter describes how to manage redundancy in line modules, switch route processor (SRP) modules, switch fabric modules (SFMs), I/O modules, and I/O adapters (IOAs) in E Series routers.

This chapter contains the following sections:

- Line Module Redundancy Overview on page 7
- Monitoring Line Module and SRP Module Redundancy on page 19
- Managing Port Redundancy on page 23

Line Module Redundancy Overview

You can install an extra line module in a group of identical line modules to provide redundancy if one of the modules fails.

The process by which the router switches to the spare line module is called *switchover*. During switchover, the line, circuit, and IP interfaces on the I/O module or one or more IOAs appear to go down temporarily. The duration of the downtime depends on the number of interfaces and the size of the routing table, because the router must reload the interface configuration and the routing table from the SRP module.

If the line module software is not compatible with the running SRP module software release, a warning message appears on the console.

Module Requirements

The requirements for line module redundancy depend on the type of router that you have.



NOTE: The information in this section does not apply to the ERX310 Broadband Services Router, which does not support line module redundancy.

ERX7xx Models and ERX14xx Models

To use this feature on ERX7xx models and ERX14xx models, you must also install a redundancy midplane and a redundancy I/O module. For a detailed explanation

of how the router provides redundancy for line modules and procedures for installing the modules, see the *ERX Hardware Guide*.

E120 and E320 Routers

To configure line module redundancy on the E120 or E320 Broadband Services router, you must also install an ES2-S1 Redund IOA in either slot 0 or slot 11. The ES2-S1 Redund IOA is a full-height IOA. For a detailed explanation of how the router provides redundancy for line modules and procedures for installing the modules, see the *E120 and E320 Hardware Guide*.

On E120 and E320 routers, each side of the chassis is treated as a redundancy group. The lowest numbered slot for each side acts as the spare line module, providing backup functionality when an ES2-S1 Redund IOA is located directly behind it. When the line module does not contain an ES2-S1 Redund IOA, it is considered a primary line module.

The router accepts the following redundancy groups:

- ES2 4G LM as backup and ES2 4G LM as primary
- ES2 10G Uplink LM and ES2 10G Uplink LM as primary
- ES2 10G LM as backup and ES2 10G LM
- ES2 10G ADV LM as backup and ES2 10G ADV LM as primary
- ES2 10G ADV LM as backup and ES2 10G LM as primary

Also, you cannot configure redundancy for the ES2-S1 Service IOA.

IOA Behavior When the Router Reboots

On E120 and E320 routers, switchover is based on the combined states of the line module and the IOAs that are installed in the affected slot.

When the router reboots and the formerly configured primary line module is not present, or is present and fails diagnostics, it switches to a spare line module and takes inventory of the IOAs. If the IOA is present and new, the router reverts back to the primary line module so that the spare line module can service other active primary line modules.

When the router reboots and there is a slot that contains a line module and one active and one inactive IOA, the inactive IOA remains in that state.

Line Module Behavior When Disabling or Enabling IOAs

On E120 and E320 routers, a line module reboots when you issue the **adapter disable** or **adapter enable** commands for an associated IOA.

When you issue the **adapter disable** or **adapter enable** commands, the line module (primary or spare) currently associated with that IOA reboots. If the IOA is protected by a line module redundancy group, an automatic line module redundancy switchover or revert can be triggered by the line module reboot. To prevent undesired line

module redundancy actions, issue the **redundancy lockout** command for the primary line module slot before issuing the **adapter disable** or **adapter enable** commands.

Automatic Switchover

Provided you have not issued the **redundancy lockout** command for the primary line module, the router switches over to the spare line module automatically if it detects any of the following failures on the primary line module:

- Power-on self-test (POST) failure
- Software-detected unrecoverable error
- Software watchdog timer expiration
- Primary line module failure to respond to keepalive polling from the SRP module
- Removal, disabling, or reloading of the primary line module
- Missing or disabled primary line modules when the router reboots
- Resetting the primary line module using the concealed push button

Limitations of Automatic Switchover

If automatic switchover is enabled on a slot (the default configuration) and a spare line module is available, issuing some CLI commands for the primary line module causes a switchover (Table 3 on page 9).

You can also disable automatic switchover on individual slots. For more information, see “Configuring Line Module Redundancy” on page 10.

Table 3: Commands That Can Cause Automatic Switchover

Command	Reason for Automatic Switchover
slot disable <i>primary-line-module-slot</i>	The command disables the primary line module but not the primary I/O module or IOAs.
reload slot <i>primary-line-module-slot</i>	The command is equivalent to pushing the reset button on the primary line module.

Reversion After Switchover

You can install only one spare line module in the group of slots covered by the redundancy midplane or redundancy group. If the router is using the spare line module, no redundancy is available. Reverting to the primary module as soon as possible is desirable. By default, the router does not automatically revert to the primary module after switchover; however, you can configure it to do so. (See “Configuring Line Module Redundancy” on page 10.) Before reversion can take place, the primary line module must complete the POST diagnostics.

Configuring Line Module Redundancy

You can modify the default redundancy operations on the router as follows:

- Disable automatic switchover on a slot.
- Enable automatic reversion after switchover.

redundancy lockout

- Use to prevent the router from switching automatically to a spare line module if the primary module in the specified slot fails.
- The **redundancy force-switchover** command overrides the **redundancy lockout** command.
- Example

```
host1(config)#redundancy lockout 5
```

- Use the **no** version to restart redundancy protection for the slot.
- See redundancy lockout.

redundancy revertive

- Use to enable the router to revert from all spare line modules to the associated primary line modules automatically.
- Reversion takes place when the primary line module is again available unless you specify a time of day. In that case, reversion takes place only when the primary module is available and after the specified time.
- Example

```
host1(config)#redundancy revertive 23:00:00
```

- Use the **no** version to disable automatic reversion.
- See redundancy revertive.

Managing Line Module Redundancy

When the router is running and a redundancy group is installed, you can manage the redundancy situation as follows:

- Force switchover manually.
- Force reversion manually.

redundancy force-switchover

- Use to force the router to switch from the primary line module in the specified slot or the primary SRP module to the spare line module or SRP module.
- The command causes a single switchover. When you reboot, the router reverts to the configured setting for this slot.
- The **redundancy force-switchover** command overrides the **redundancy lockout** command.
- Example

```
host1#redundancy force-switchover 5
```

- There is no **no** version.
- See redundancy force-switchover.

redundancy revert

- Use to force the router to revert to the primary line module in the specified slot.
- The router acts on this command immediately unless you specify a time or a time and date that the action is to take place.
- The command causes a single reversion. When you reboot, the router uses the configured setting for this slot.
- Example

```
host1#redundancy revert 4 23:00:00 5 September 200X
```

- There is no **no** version.
- See redundancy revert.

SRP Module Redundancy

This section covers general issues of SRP module redundancy. It does not cover NVS cards or the behavior on systems running high availability features such as hitless SRP switchover. For information about managing NVS in a router that contains two SRP modules, see *Managing Flash Cards on SRP Modules* in the *JUNOS System Basics Configuration Guide*. For information about managing high availability in a router, see “Managing Module Redundancy” on page 7.

The information in this section does not apply to the ERX310 router, which does not support SRP module redundancy. For this reason, any references to synchronization that may appear in command output or system messages do not apply to the ERX310 router.

SRP Module Behavior

The SRP module uses a 1:1 redundancy scheme. When two SRP modules are installed in the router, one acts as a primary and the second as a redundant module. On ERX7xx models, ERX14xx models, and the ERX310 router, both SRP modules share a single SRP I/O module located in the rear of the chassis. On the E120 and E320 routers, both SRP modules share an SRP IOA located in the rear of the chassis.

After you install two SRP modules, the modules negotiate for the primary role. A number of factors determine which module becomes the primary; however, preference is given to the module in the lower slot. The SRP modules record their latest roles and retain them the next time you switch on the router.

With the default software settings, if the primary SRP module fails, the redundant SRP module assumes control without rebooting itself. For information about preventing the redundant SRP module from assuming control, see “Managing SRP Module Redundancy” on page 16.

On E120 and E320 routers, the switch fabric is distributed between the SFMs and the SRP modules. If the primary SRP module fails a diagnostic test on its resident slice of switch fabric, then it abdicates control to the redundant SRP module if both of the following are true:

- The standby SRP module does not indicate any error.
- The standby SRP module passes diagnostics on its attached fabric slice.

When the redundant SRP module assumes control, the following sequence of events occurs:

1. The original primary SRP module reboots and assumes the redundant role.
2. The redundant SRP module restarts and assumes the primary role without reloading new code. (When upgrading software, you must reload the software on the redundant SRP module. See *Installing JUNOS Software* in the *JUNOS System Basics Configuration Guide*.)
3. All line modules reboot.

The following actions activate the redundant SRP module:

- Failure of the primary SRP module (hardware or software)
- Pushing the recessed reset button on the primary SRP module (See Figure 1 on page 13 and Figure 2 on page 14.)
- Issuing the **srp switch** command
- Issuing the **redundancy force-switchover** command

Figure 1: SRP Module on ERX7xx Models and ERX14xx Models





Figure 2: SRP Module on the E120 and E320 Routers

Specifying the Configuration for Redundant SRP Modules

On a router with redundant SRP modules, you can specify the configuration that both the primary and redundant modules load in the event of a reload or switchover. A switchover can result from an error on the primary SRP module or from an **srp switch** command. The following behavior takes place only in the event of a cold restart; it does not take place in the event of a warm restart.

When you arm a configuration (.cnf) file by issuing the **boot config cnfFilename** command, a subsequent SRP switchover causes the redundant SRP module to take the role of primary SRP module with the configuration specified by the .cnf file. The new primary SRP module does not use the running configuration.

If you want the redundant SRP module to instead use the running configuration when it takes the primary role, then you must first arm a configuration file with the **boot config cnfFilename once** command. To exhaust the **once** option, you must then cause the redundant SRP module to reload for some reason, such as by issuing a **reload** command or by arming a new JUNOS software release or a hotfix file.

When a switchover subsequently occurs, the redundant SRP module reloads with the running configuration and takes the primary role. For more information about the **boot config** command, see *Booting the System* in the *JUNOS System Basics Configuration Guide*.

Installing a Redundant SRP Module

You can install a redundant SRP module into a running router, provided that the redundant SRP module has a valid, armed software release on its NVS card. Access to a software release in NVS ensures that the redundant SRP module can boot; the release need not be the same as that on the primary SRP module.



WARNING: Do not insert any metal object, such as a screwdriver, or place your hand into an open slot or the backplane when the router is on. Remove jewelry (including rings, necklaces, and watches) before working on equipment that is connected to power lines. These actions prevent electric shock and serious burns.



CAUTION: When handling modules, use an antistatic wrist strap connected to the router's ESD grounding jack, and hold modules by their edges. Do not touch the components, pins, leads, or solder connections. These actions help to protect modules from damage by electrostatic discharge.

To install a redundant SRP module into a running router, follow these steps:

1. Install the redundant SRP module into the open SRP slot (slot 6 or 7 for ERX14xx models, the E120 router, and the E320 router; slot 0 or 1 for ERX7xx models).

For detailed information about installing the SRP module, see the *ERX Hardware Guide* or the *E120 and E320 Hardware Guide*.

2. Wait for the redundant SRP module to boot, initialize, and reach the standby state.

When the module is in standby state, the REDUNDANT LED is on and the ONLINE LED is off. If you issue the **show version** command, the state field for the slot that contains the redundant SRP module is standby.

3. Synchronize the NVS file system of the redundant SRP module to that of the primary SRP module.
-



NOTE: The SRP module reboots after synchronization is complete.

reload slot

- Use to reboot a selected slot on the router.
- If you specify a slot on the E120 or E320 router that contains an SRP module, you reboot the SC subsystem on that slot by default. You do not, however, reboot the fabric slice that resides on the slot.

- Use the **srp** keyword to reboot the portion of the SC subsystem that resides on a specified SRP module.
- Use the **fabric** keyword to reboot the fabric slice that resides on the specified SRP module.
- Example 1—Reboots the module in slot 7
`host1#reload slot 7`
- Example 2—Reboots the SC on the SRP module in slot 7 (applies only to E120 and E320 routers)
`host1#reload slot 7 srp`
- There is no **no** version.
- See reload slot.

synchronize

- Use to force the file system of the redundant SRP module to synchronize with the NVS file system of the primary SRP module.
- If you synchronize the redundant SRP module with the primary SRP module and the redundant module is armed with a release different from the one it is currently running, the redundant SRP module is automatically rebooted to load the armed release.
- Optionally, you can use the **low-level-check** keyword to force the router to validate all files or only configuration files in NVS, and to synchronize all files that failed the checksum test during the **flash-disk-compare** command as well as any other files that are unsynchronized. See *Managing Flash Cards on SRP Modules* in the *JUNOS System Basics Configuration Guide* for details.
- Examples
`host1#synchronize`
`host1#synchronize low-level-check all`
`host1#synchronize low-level-check configuration`
- There is no **no** version.
- See synchronize.

Managing SRP Module Redundancy

You can prevent the redundant SRP module from taking over when:

- The primary SRP module experiences a software failure.
- You push the reset button on the primary SRP module.



NOTE: If you do not configure this option, when troubleshooting an SRP module, disconnect the other SRP module from the router. This action prevents the redundant SRP module from taking over if you push the reset button on the primary SRP module.

To configure this option:

1. Issue the **disable-switch-on-error** command.
2. Synchronize the NVS file system of the redundant SRP module to that of the primary SRP module.

disable-switch-on-error

- Use to prevent the redundant SRP module from taking over if the primary SRP module experiences a software failure or if you push the reset button on the primary SRP module.
- Issue the **synchronize** command immediately before you issue this command.
- If you issue the **disable-switch-on-error** command, and later issue the **srp switch** command, the redundant SRP module waits about 30 seconds before it takes over from the primary SRP module.
- Example

```
host1(config)#disable-switch-on-error
```

- Use the **no** version to revert to the default situation, in which the redundant SRP module takes over if the primary SRP module experiences a software failure.
- See disable-switch-on-error.

synchronize

- Use to force the NVS file system of the redundant SRP module to synchronize with the NVS file system of the primary SRP module.
- If you synchronize the redundant SRP module with the primary SRP module and the redundant module is armed with a release different from the one it is currently running, the redundant SRP module is automatically rebooted to load the armed release.
- Optionally, you can use the **low-level-check** keyword to force the router to validate all files or only configuration files in NVS, and to synchronize all files that failed the checksum test during the **flash-disk-compare** command as well as any other files that are unsynchronized. See *Managing Flash Cards on SRP Modules* in the *JUNOS System Basics Configuration Guide* for details.
- Examples

```
host1#synchronize
host1#synchronize low-level-check all
host1#synchronize low-level-check configuration
```

- There is no **no** version.
- See synchronize.

Switching to the Redundant SRP Module

To switch immediately from the primary SRP module to the redundant SRP module, issue the **redundancy force-switchover** command or the **srp switch** command. You can configure the router to prompt you if the modules are in a state that could lead to loss of configuration data or NVS corruption.

redundancy force-switchover

- Use to force the router to switch from the primary line module in the specified slot or the primary SRP module to the spare line module or SRP module.
- The command causes a single switchover. When you reboot, the router reverts to the configured setting for this slot.
- With the **srp** option, the command is equivalent to the **srp switch** command.
- The **redundancy force-switchover** command overrides the **redundancy lockout** command.
- Example

```
host1#redundancy force-switchover 5
```

- There is no **no** version.
- See redundancy force-switchover.

srp switch

- Use to switch from the primary SRP module to the redundant SRP module.
- When the high availability state is active, this command delays until all transaction data, up to when you issue the command, has been mirrored to the standby SRP module. This preserves legacy behavior requiring that SRP modules be synchronized before the switchover.
- If you specify the **force** keyword, the procedure fails if the SRP modules are in certain states, such as during a synchronization. In these cases, the router displays a message that indicates that the procedure cannot currently be performed and the reason why. However, if the SRP modules are in other states that could lead to a loss of configuration data or an NVS corruption, the router displays a message that explains the state of the SRP modules, and prompts you to confirm (enter yes or no) whether you want to proceed.
- If you do not specify the **force** keyword, the procedure fails if the SRP modules are in any state that could lead to a loss of configuration data or an NVS corruption, and the router displays a message explaining the command failure.
- When you issue this command, the router prompts you for a confirmation before the command takes effect.
- If you issue the **disable-switch-on-error** command and later issue the **srp switch** command, the redundant SRP module waits about 30 seconds before it takes over from the primary SRP module.

- If the router does not contain a redundant SRP module, this command has no effect.
- Example

```
host1#srp switch
host1#srp switch force
```
- There is no **no** version.
- See srp switch.

Upgrading Software on a Redundant SRP Module

For information about upgrading software on SRP modules on ERX7xx models, ERX14xx models, or the ERX310 router, see *Installing JUNOS Software* in the *JUNOS System Basics Configuration Guide*.

Monitoring the Status LEDs

You can determine the redundancy state of line modules and SRP modules by examining their status LEDs. See Table 4 on page 19 for a description of the LEDs functions. In addition, if you issue the **show version** command, the state field for the slot that contains the redundant SRP module indicates standby.

Table 4: Function of the Online and Redundant LEDs

Online LED	Redundant LED	State of the Module
Off	Off	Module is booting or is an inactive primary line module.
On	Off	Module is active, but no redundant module is available.
Off	On	Module is in standby state.
On	On	Module is active, and a redundant module is available.

Monitoring Line Module and SRP Module Redundancy

You can use **show** commands to monitor the status of redundancy groups, line modules, and SRP modules.



NOTE: For more information about monitoring high availability, see “Managing Module Redundancy” on page 7.

show environment

- Use to display information about the hardware installed for redundancy.

- See *Managing the System* in the *JUNOS System Basics Configuration Guide*, for details and examples.
- See show environment.

show hardware

- Use to display detailed information about the line modules and SRP modules.
- See *Monitoring Modules* in the *JUNOS System Basics Configuration Guide* for details and examples.
- See show hardware.

show redundancy

- Use to display the configuration for line module redundancy and SRP redundancy.
- Field descriptions
 - SRP
 - high-availability state—State of the high availability mode (disabled, active, or pending)
 - current redundancy mode—Redundancy mode currently being used by this router (high-availability or file-system-synchronization)
 - last activation type—Last type of activation that occurred on this router (that is, the method by which the SRP last booted [cold-start or warm-start])
 - Criteria Preventing High Availability from being Active—Criteria required for HA to be active.
 - slot—Slot in which the line module resides
 - hardware role—Function of the line module: primary or spare
 - lockout config—Status of redundancy on this line module
 - protected—Line module redundancy is enabled
 - locked out—Line module redundancy is disabled
 - backed up by slot—Slot that contains the line module that is a spare for this primary line module
 - sparing for slot—Slot that contains the primary line module for which this line module is a spare
 - revert at—Time at which you want the line module to revert
 - midplane type—Identifier for the type of midplane
 - midplane rev—Hardware revision number of the redundancy midplane
 - fabric slice redundancy—Status of the fabric slice on the SRP modules or SFMs on the E120 and E320 routers
 - slot—Slot in which the fabric slice resides

- state—State of the fabric slice (online, not present)
- type—Identifier for the type of hardware (SRP module or SFM)

■ Example 1

In the following example, the user issues a **show redundancy** command, and then a **redundancy force switchover** command. Finally, the user issues the **show redundancy line-card** command to display information specific to the line modules. The two displays show how the states of the primary and spare line modules change.

```
host1#show redundancy

SRP
---

high-availability state: disabled
current redundancy mode: file-system-synchronization
last activation type: cold-start

Criteria Preventing High Availability from being Active
-----
              criterion                      met
-----
High Availability mode configured?         No
Mirroring Subsystem present?               No

Line Card
-----

automatic reverting is off

              backed
              up
              by
slot  hardware  lockout  slot  sparing  revert
-----
0      spare    ---      ---      ---      ---
2      primary  protected ---      ---      ---
12     ---      ---      ---      ---      ---

              midplane  midplane
slots         type      rev
-----
0 - 5         6         0
```

```
host1#redundancy force-switchover 2
host1#show redundancy line-card
```

```
automatic reverting is off

              backed
              up
              by
slot  hardware  lockout  slot  sparing  revert
-----
0      spare    ---      ---      2        ---
2      primary  protected  0      ---      ---
```

```

12      ---      ---      ---      ---      ---

slots      midplane  midplane
           type      rev
-----
0 - 5      6         0

```

- Example 2—Displays the redundancy status on an E320 router

```
host1#show redundancy
```

```
SRP
```

```
---
```

```

high-availability state: active
current redundancy mode: high-availability
last activation type:    cold-start

```

```
Line Card
```

```
-----
```

```
automatic reverting is off
```

slot	hardware role	lockout config	backed up by slot	sparing for slot	revert at
0	spare	---	---	---	---
2	primary	protected	---	---	---
4	primary	protected	---	---	---

```
fabric slice redundancy: none
```

slot	state	type
6	online	SFM-100
7	online	SFM-100
8	---	---
9	---	---
10	---	---

- See show redundancy.

show version

- Use to display information about each module in the router.

See *Managing the System* in the *JUNOS System Basics Configuration Guide*, for details and examples.

- See show version.

Managing Port Redundancy

For information on port redundancy, see the *JUNOS Physical Layer Configuration Guide*. For information on managing port redundancy on Gigabit Ethernet I/O modules, see *Managing Port Redundancy on Gigabit Ethernet I/O Modules* in the *JUNOS Physical Layer Configuration Guide*.

For information on redundancy and interface distribution of tunnel-service interfaces see *Redundancy and Interface Distribution of Tunnel-Service Interfaces* in the *JUNOS Physical Layer Configuration Guide*.

Chapter 3

Managing Stateful SRP Switchover

This chapter describes how to manage Juniper Networks Stateful SRP Switchover (also referred to as high availability or HA) software features for the E Series router. Use this chapter with “Managing Module Redundancy” on page 7 to fully manage the SRP features.

This chapter contains the following sections:

- Understanding Stateful SRP Switchover on page 25
- Platform Considerations on page 26
- Redundancy Modes of Operation on page 26
- Understanding SRP State Behavior on page 28
- Application Support on page 31
- Before Activating High Availability on page 40
- Activating High Availability on page 40
- Deactivating High Availability on page 41
- Upgrading Software on page 42
- Monitoring Stateful SRP Switchover on page 43

Understanding Stateful SRP Switchover

Stateful SRP switchover is the idea of reducing or eliminating single points of failure. When applied to the E Series router, stateful SRP switchover provides both hardware-specific and software-specific methods to ensure minimal downtime and ultimately improve the performance of your network.

For hardware components, Juniper Networks provides redundancy solutions to ensure that the router continues to operate in the event of a hardware fault. This redundancy can exist on various router models in the form of multiple power supplies, cooling fans, switching planes, routing engines and, in some cases, interfaces. Redundancy also allows for hot-swapping various components within your Juniper Networks router.



NOTE: For information about E Series hardware redundancy features, see the *ERX Hardware Guide* or the *E120 and E320 Hardware Guide*.

Platform Considerations

Stateful SRP switchover is supported on all E Series routers except for the ERX310 Broadband Services Router.

For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models and ERX14xx models.
- See the *E120 and E320 Module Guide* for modules supported on the E120 and E320 Broadband Services Routers.

Module Requirements

The following table lists which SRPs support or do not support the high availability mode (stateful SRP switchover) feature.

SRP Model	Supported
SRP-5G	No
SRP-5G +	Yes
SRP-10G	Yes
SRP-40G	No
SRP-40G PLUS	Yes
SRP-100	Yes



NOTE: Stateful SRP switchover requires two SRP modules with 1 GB of memory or more.

Redundancy Modes of Operation

The switch route processor (SRP) modules can operate in one of two redundancy modes—file system synchronization and high availability.

File System Synchronization Mode

File system synchronization is the default behavior mode for E Series routers that contain redundant SRPs. Available only to SRP modules, this mode has been available since the JUNOS Software 2.x release. In this mode:

- Files and data (for example, configuration files and releases) in nonvolatile storage (NVS) remain synchronized between the primary and standby SRP modules.
- SRP modules reload all line modules and restart from saved configuration files.
- If the active SRP module switches over to the standby SRP, the router cold-restarts as follows:
 - All line modules are reloaded.
 - User connections are lost, and forwarding through the chassis stops until the router SRP module recovers.
 - The standby SRP module boots from the last known good configuration from NVS.

For additional information about the default SRP functionality, see “Managing Module Redundancy” on page 7.

High Availability Mode

Currently applicable to the SRP module, the Juniper Networks high availability mode uses an initial bulk file transfer and subsequent, transaction-based mirroring to ensure rapid SRP module recovery after a switchover. This process is referred to in this chapter as *stateful SRP switchover*.

In addition to keeping the contents of NVS, high availability mode keeps state and dynamic configuration data from the SRP memory synchronized between the primary and standby SRP modules.

When stateful SRP switchover is enabled, an SRP switchover keeps line modules up and forwarding data, and the newly active SRP module continues from the point of switchover.

By using transaction-based mirroring instead of file synchronization, high availability mode keeps the standby SRP module synchronized with the active SRP module. Mirroring occurs from memory on the active SRP module to memory on the standby SRP module by way of transactions. When a transaction is committed on the active SRP module, the data associated with the transaction is sent to the standby SRP module.

In high availability mode:

- The contents of the NVS in the primary and standby SRP modules remain synchronized.



NOTE: Configuration files are always synchronized. Nonconfiguration files are synchronized when the **disable-autosync** command has not been configured; this is the default case. When the **disable-autosync** command has been configured, nonconfiguration files are not synchronized.

- If a switchover occurs:

- The standby SRP module warm-restarts using the mirrored data to restore itself to the state of the system before the switchover.
- During the warm restart:
 - User connections remain active, and forwarding continues through the chassis.
 - New user connection attempts during switchover are denied until switchover is complete.
 - New configuration changes are prevented until switchover is complete (or after 5 minutes).

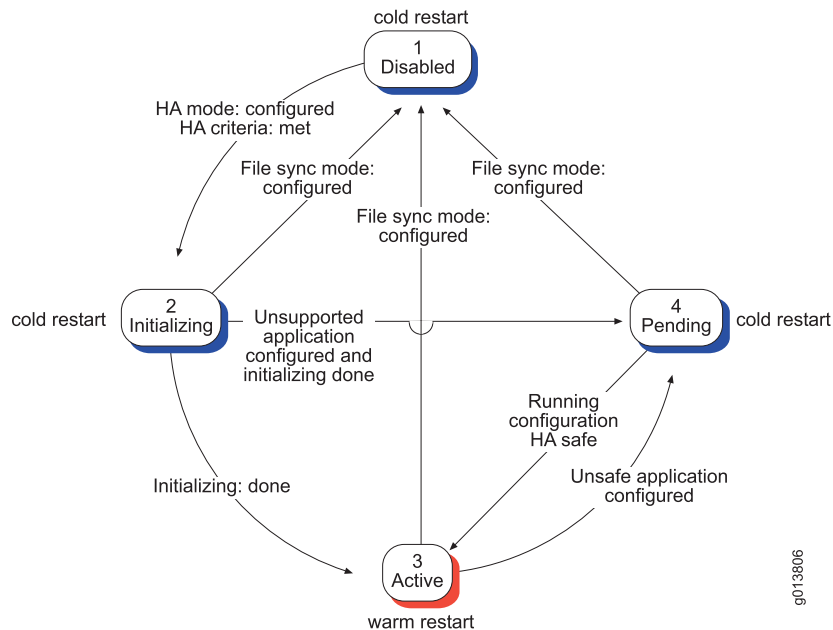


NOTE: If the switchover does not finish within 5 minutes, the SRP module cancels the operation and reenables CLI configuration.

Understanding SRP State Behavior

The SRP progresses through various high availability states. These states are illustrated in Figure 3 on page 28.

Figure 3: High Availability States



g013806

Disabled State

The initial, default state for high availability mode is disabled. While in this state, the router continues to use file system synchronization. If a switchover occurs while the router is in this state, the standby SRP module performs a cold restart.

The router enters this state when you power up the router or when the router warm-restarts from an SRP switchover.

After you enable high availability, the system must meet the following criteria before it can enter the initializing state:

- High availability mode is configured.
- Active SRP hardware supports high availability.
- Network core dump feature is disabled.
- Running configuration allows high availability to operate (that is, no unsupported applications are configured).
- Standby SRP hardware supports high availability.
- Standby SRP module is online and capable of mirroring.
- Standby SRP module is running the same release.

During the disabled state:

- If any one criterion is not met, the system remains in the disabled state, until the criterion is met.
- If a switchover occurs while the system is in the disabled state, the system cold-restarts.

While in the disabled state, the system operates as if it were configured for file system synchronization (for example, NVS is synchronized every 5 minutes, if autosynchronization is enabled).

If all criteria are met, high availability mode transitions to the initialization state.

Initializing State

After the SRP module transitions into the initializing state, bulk synchronization of the memory and NVS occurs. This includes the following:

- File synchronization of the primary NVS with the standby NVS
- Mirroring of appropriate state and dynamic configuration information from the active SRP (memory) to the standby SRP (memory)



NOTE: Depending on the size of the configuration, this process can take several minutes.

During the initializing state:

- If an unsupported application is configured during initialization, the system completes initializing and enters the pending state.
- If any other criterion becomes false (or is no longer met), the system enters the disabled state.
- If a switchover occurs while the system is in this state, the system cold-restarts.

After initialization is completed, the system enters the active state.

Active State

During the active state, the data that was synchronized from the active SRP module to the standby SRP module during initialization remains synchronized through mirroring updates.

Mirroring updates occur as follows:

1. When making changes or updates, applications create individual transactions, perform the updates on the active SRP module, and post the transactions.
2. Following the updates, the active SRP module sends the changes to the standby SRP module.
3. The standby SRP module replays the updates (in the order in which they were committed on the active SRP module) and makes the appropriate changes for each changed application.
4. Updates that need to be stored in NVS (that is, for static configurations) are updated in NVS.



NOTE: While in the active and pending states, the CLI **synchronize** command does not update configuration files; these files are updated by the mirroring process.

During the active state:

- If a switchover occurs while the router is in the active state, the standby SRP module performs a warm restart (that is, stateful SRP switchover is in effect); the standby SRP module uses the configuration located in NVS.
- If an unsupported application is configured, the system transitions to the pending state.
- If any other criterion changes (is no longer met), the system transitions to the disabled state.



NOTE: Changes made in manual commit mode are maintained, uncommitted, in the standby SRP memory until a trigger to commit occurs; if a switchover occurs while in this mode, the standby SRP module uses the configuration in memory.

Pending State

The system transitions to the pending state if an unsupported application is configured. When a transition to the pending state occurs, the system generates SNMP traps and log messages.

How the router behaves depends on which HA state the application is in when it shifts to a pending state:

- From disabled state—The router remains in the disabled state.
- From initializing state—The router completes the initializing state and transitions to the pending state after initialization is complete.
- Active State—The router transitions to the pending state.

The system remains in the pending state until the configuration of the unsupported application is removed. However, even though it is in the pending state, the system continues mirroring updates from the primary SRP module to the standby SRP module.



NOTE: You can use the **show redundancy srp** command to display the name of any unsupported applications that are configured.

If a switchover occurs while the system is in the pending state, the system cold-restarts.

Application Support

Applications are either supported or unsupported by stateful SRP switchover.

- Supported—You can configure supported applications without having any adverse impact to stateful SRP switchover. When a switchover occurs, supported applications can react to switchovers in one of two different ways:
 - Gracefully recover using mirrored static and dynamic information (for example, IP, PPP, and PPPoE)
 - Recover using static configuration only; that is, no runtime state is restored after a switchover. Dynamic configuration and state information are lost. (For example, CLI sessions are restarted, telnet sessions are dropped, multicast routes must be rebuilt, and so on.)
- Unsupported—We recommend that you not configure unsupported applications on a chassis running in high availability mode. Although configured unsupported applications suspend high availability or prevent high availability from becoming active, they do not cause any problems with the function of the router.

Table 5 on page 32 indicates which applications support or do not support stateful SRP switchover.

Table 5: Application Support for Stateful SRP Switchover

Application	Supported	Unsupported	Notes
Physical Layer Protocols			
DS1	✓	–	–
DS3	✓	–	–
HDLC	✓	–	–
SONET/SDH	✓	–	–
SONET/SDH VT	✓	–	–
Link-Layer Protocols			
ATM	✓	–	Static and dynamic interfaces, with the exception of ATM subscribers, are supported. In this case, <i>ATM subscribers</i> refers to a technology on the E Series router where the ATM layer does authentication (that is, not PPP or IP subscriber manager).
ATM 1483 bulk configuration of dynamic interfaces	✓	–	–
Bridged Ethernet	✓	–	–
Cisco HDLC	✓	–	–
Ethernet (with and without VLANs)	✓	–	–
Frame Relay	✓	–	–
PPP	✓	–	–
PPPoE	✓	–	–
Transparent bridging	✓	–	–
Unicast Routing			
Access Routes	✓	–	–

Table 5: Application Support for Stateful SRP Switchover *(continued)*

Application	Supported	Unsupported	Notes
BFD	✓	–	During a stateful SRP switchover, the BFD transmit interval is set to 1000 ms with a detection multiplier of 3. These values result in a liveness detection interval of 3000 ms. This longer interval helps prevent a BFD timeout during the switchover. BFD negotiates the interval with the remote peer before applying the temporary change. The BFD timers revert back to the configured values after 15 minutes (the maximum duration for graceful restart completion).
BGP	✓	–	Supported for IPv4 only when the graceful restart extension is enabled. Does not support graceful restart for IPv6 address families.
FTP	✓	–	Static recovery support only.
IP	✓	–	–
IPv6	✓	–	–
IPv6 neighbor discovery	✓	–	IPv6 neighbor discovery characteristics are replayed during switchover. Line modules do not flush neighbor discovery information during the switchover.
IPSec Transport	–	✓	–
IPSec Tunnels	✓	–	Completed IKE phase 1 and phase 2 negotiations supported only.
IS-IS	✓	–	Supported only when the graceful restart extension is enabled.
IS-ISv6	✓	–	Supported only when the graceful restart extension is enabled.
OSPFv2		–	Supported only when the graceful restart extension is enabled.
OSPFv3	✓	–	Supported only when the graceful restart extension is enabled.
RIP	✓	–	Static recovery support only.

Table 5: Application Support for Stateful SRP Switchover *(continued)*

Application	Supported	Unsupported	Notes
Static Routes (IP and IPv6)	✓	–	After all high-priority interfaces have been replayed from NVS and mirrored storage, static routes are replayed from NVS, followed by replay of low-priority interfaces from NVS and mirrored storage. This behavior enables static routes that are dependent on high-priority interfaces to be resolved quickly and installed in the IP routing table.
Telnet	✓	–	Static recovery support only.
IPv4 Multicast Routing			
Multicast Routing	✓	–	Static recovery support only. During switchover, the system mirrors the multicast queue so that IP can use the same queue without needing to re-create a different connection.
DVMRP	✓	–	Static recovery support only. DVMRP gives the restart complete indication to the IP routing table after getting a peer update (60-second timeout).
IGMP	✓	–	<p>IC IGMP deletes its interface and membership state on SRP failover (controller down). As part of SRP warm start, IGMP interfaces are reconfigured from NVS and dynamic IGMP interfaces are reconfigured from mirrored storage. IGMP hosts are queried as IP interfaces come back up, the join state is re-established, and SC IGMP state is created.</p> <p>After the maximum query response time (across all interfaces) expires to allow hosts to re-establish join state, IGMP notifies MGTM that graceful restart is complete.</p>

Table 5: Application Support for Stateful SRP Switchover *(continued)*

Application	Supported	Unsupported	Notes
MGTM	✓	–	<p>On SRP failover, old mroutes are retained on the line module to preserve multicast forwarding; cache-misses to the SRP are disabled. When MGTM warm starts on the SRP, it reads the NVS configuration and enables multicast routing. When IGMP, DVMRP, and PIM have completed graceful restart and the IP route table multicast-view has completed graceful restart, old mroutes are deleted from the line module and cache-misses to the SRP are enabled. This triggers re-creation of mroutes and establishes the current multicast forwarding state.</p> <p>Although cache-misses to the SRP module are disabled, forwarding is preserved for old multicast joins to downstream routers and hosts. However, forwarding for new multicast joins requested by downstream routers and hosts after SRP module switchover is not provided until cache-misses are re-enabled.</p>
PIM	✓	–	<p>Static recovery support only. For warm start, PIM interfaces are reconfigured from NVS. A Hello message with a new Generation ID is issued as IP interfaces come up. A neighbor that receives this Hello determines that the upstream neighbor has lost state and needs to be refreshed. A VR-global configurable graceful restart timer is required for PIM to time out the re-establishment of the join state for sparse-mode interfaces. After this timer expires, PIM notifies MGTM that graceful restart is complete.</p>
IPv6 Multicast Routing			
Multicast Routing	✓	–	<p>Static recovery support only. During switchover, the system mirrors the multicast queue so that IPv6 can use the same queue without needing to re-create a different connection.</p>

Table 5: Application Support for Stateful SRP Switchover *(continued)*

Application	Supported	Unsupported	Notes
MGTM	✓	–	<p>On SRP failover, old mroutes are retained on the line module to preserve multicast forwarding; cache-misses to the SRP are disabled. When MGTM warm starts on the SRP, it reads the NVS configuration and enables multicast routing. When MLD and PIM have completed graceful restart and the IPv6 route table multicast-view has completed graceful restart, old mroutes are deleted from the line module and cache-misses to the SRP are enabled. This triggers re-creation of mroutes and establishes the current multicast forwarding state.</p> <p>Although cache-misses to the SRP module are disabled, forwarding is preserved for old multicast joins to downstream routers and hosts. However, forwarding for new multicast joins requested by downstream routers and hosts after SRP module switchover is not provided until cache-misses are re-enabled.</p>
MLD	✓	–	<p>IC MLD deletes its interface and membership state on SRP failover (controller down). As part of SRP warm start, MLD interfaces are reconfigured from NVS and dynamic IMLD interfaces are reconfigured from mirrored storage. MLD hosts are queried as IPv6 interfaces come back, the join state is re-established, and SC MLD state is created. After the maximum query response time (across all interfaces) expires to allow hosts to re-establish join state, MLD notifies MGMTv6 that graceful restart is complete.</p>
PIM	✓	–	<p>Static recovery support only. For warm start, PIM interfaces are reconfigured from NVS and a Hello message with a new Generation ID is issued as IPv6 interfaces come up. A neighbor that receives this Hello determines that the upstream neighbor has lost state and needs to be refreshed. A VR-global configurable graceful restart timer is required for PIM to time out the re-establishment of the join state for sparse-mode interfaces. After this timer expires, PIM notifies MGMT that graceful restart is complete.</p>

Table 5: Application Support for Stateful SRP Switchover *(continued)*

Application	Supported	Unsupported	Notes
Multiprotocol Label Switching			
MPLS	✓	–	<p>MPLS is HA-unsafe during a graceful restart. It is HA-unsafe until all the configured MPLS signaling protocols have completed their graceful restart procedures and any stale forwarding elements have been flushed from the line modules.</p> <p>If you force an SRP switchover while MPLS is HA-unsafe, the SRP module switches but the SRP module and the line modules undergo a cold restart.</p> <p>If the primary SRP module resets while MPLS is HA-unsafe, the router undergoes a cold restart.</p> <p>MPLS over IPv6 supports HA. This functionality enables BGP to support graceful restart for IPv6 labeled addresses.</p>
BGP signaling	✓	–	—
LDP signaling	✓	–	<p>To provide uninterrupted service during an SRP switchover in a scaled configuration, such as one with 32,000 Martini circuits, set the LDP graceful restart reconnect time to the maximum 300 seconds and set the LDP graceful restart recovery timer to the maximum 600 seconds. This requirement is true for all SRP switchovers, including those in the context of a unified in-service software upgrade.</p> <p>LDP signaling does not support HA for IPv6.</p>
RSVP signaling	✓	–	—
Local cross-connects between layer 2 interfaces using MPLS	✓	–	–
Policies and QoS			
Policies	✓	–	–
QoS	✓	–	Static recovery support only.
Remote Access			

Table 5: Application Support for Stateful SRP Switchover *(continued)*

Application	Supported	Unsupported	Notes
AAA	✓	–	–
DHCP External Server and Packet Trigger	✓	–	Following a switchover, the DHCP lease (that is, time remaining) is recalculated based on when the lease started. When the release timer for a client expires, the client is deleted and the access route is removed, along with the dynamic subscriber interface if it was created. If the client requests a new lease, DHCP external server resynchronizes with the new lease time.
DHCP Packet Capture	✓	–	–
DHCP Proxy Client	–	✓	–
DHCP Relay Proxy	–	–	–
DHCP Relay Server	✓	–	<p>Before HA support, clients identified by the DHCP relay server were maintained on a switchover (their state was stored to NVS); DHCP relay server always had some level of HA support.</p> <p>Currently, following a switchover, the DHCP lease (that is, time remaining) is reset. When the release timer for a client expires, the client requests a new lease. The E Series router DHCP relay server then synchronizes with the new state.</p>
DHCPv4 Local Server	✓	–	–
DHCPv6 Local Server	✓	–	<p>DHCPv6 is stateful SRP switchover (high availability) safe.</p> <p>After the router receives the first renew or rebind request from the client, it re-creates the DHCPv6 bindings using mirrored AAAA information. It does not use the mirrored DHCPv6 information as it does for other applications that support stateful SRP switchover.</p> <p>The router also re-creates the client binding for any DHCPv6 client that was assigned IPV6 prefixes from RADIUS before an SRP switchover.</p>
L2TP	✓	–	–
L2TP Dialout	–	✓	–

Table 5: Application Support for Stateful SRP Switchover (continued)

Application	Supported	Unsupported	Notes
Local Address Pools	✓	–	The internal local address server state supports only static recovery. However, the AAA application reallocates active addresses on a switchover. The resulting effect is the local address server having full HA support.
RADIUS Client	✓	–	Similar to local address server, AAA recovers disrupted RADIUS communication on a switchover. The resulting effect is the RADIUS client having full HA support.
RADIUS Dynamic-Request Server	✓	–	Static recovery support only.
RADIUS Initiated Disconnect	✓		–
RADIUS Relay Server	✓	–	–
RADIUS Route-Download Server	✓	–	–
Service Manager	✓	–	–
SRC Client	✓	–	–
TACACS +	✓	–	Static recovery support only.
Miscellaneous			
DNS	✓	–	—
DNSv6	–	✓	If DNSv6 is configured, no warning or error is displayed during a warm start. DNSv6 is subsequently configured from NVS as it is after a cold reboot.
J-Flow (IP flow statistics)	✓	–	–
Line Module Redundancy	✓	–	–
Network Address Translation	✓	–	–
NTP	✓	–	–
Resource Threshold Monitor	✓	–	–

Table 5: Application Support for Stateful SRP Switchover *(continued)*

Application	Supported	Unsupported	Notes
Response Time Reporter	✓	–	–
Route Policy	✓	–	Static recovery support only.
Subscriber Interfaces	✓	–	IPv4 only. Subscriber interfaces are not applicable to IPv6
Tunnels (GRE and DVMRP)	✓	–	–
VRRP	✓	–	Static recovery support only.



CAUTION: When IP tunnels are configured on an HA-enabled router and the Service Module (SM) carrying these tunnels is reloaded, HA transitions to the pending state. HA remains in the pending state for 5 minutes after the successful reloading of the SM. This amount of time allows for IP tunnel relocation and for the tunnels to become operational again on the SM. If an SRP switchover occurs while HA is in the pending state, the router performs a cold restart.

Before Activating High Availability

Before you activate high availability on the SRP modules, review “Managing Module Redundancy” on page 7 and any high availability–related changes to SRP management commands.

Activating High Availability

You activate high availability (stateful SRP switchover) by launching Redundancy Configuration mode and issuing the **mode high-availability** command.

When activating high availability, keep the following in mind:

- In an E Series router that supports stateful SRP switchover, both SRP modules must be running the same software release version in order to activate high availability mode.
- If high availability mode cannot become active because of different releases on the active and standby SRP modules, the system reverts to its default mode (file system synchronization).
- When active or pending, the router configuration files are mirrored from the active SRP module to the standby SRP module. All other files shared between the active and standby SRP modules are automatically synchronized using legacy synchronization methods.

To enable high availability, enter the following commands:

```
host1(config)#redundancy
host1(config-redundancy)#mode high-availability
```

mode

- Use the redundancy **mode** command to enable high availability.
- The **high-availability** keyword enables high availability mode for stateful SRP switchover. In this mode, the router uses mirroring to keep the configuration and state of the standby SRP module coordinated with the configuration and state of the active SRP module.



NOTE: High availability mode is currently available only on ERX1440, ERX1400, and ERX700 Broadband Services Routers that support dual SRPs.

- The **file-system-synchronization** keyword reverts the redundancy mode to its default. In this mode, the router uses file synchronization to keep the configuration of the standby SRP module coordinated with the configuration of the active SRP module.

- Example

```
host1(config-redundancy)#mode high-availability
```

- Use the **no** version to return high availability mode operation to its default (file system synchronization).
- See mode.

redundancy

- Use to enter Redundancy Configuration mode.
- Example

```
host1(config)#redundancy
host1(config-redundancy)#
```

- There is no **no** version.
- See redundancy.

Deactivating High Availability

To deactivate high availability support, enter the following commands:

```
host1(config)#redundancy
host1(config-redundancy)#mode file-system-synchronization
or
host1(config)#redundancy
host1(config-redundancy)#no mode
```

Setting the IP Interface Priority

During the warm restart after an SRP switchover, IP and IPv6 interfaces are replayed from NVS and from mirrored storage. High-priority IP and IPv6 interfaces are replayed first, followed by static routes, and then by low-priority IP and IPv6 interfaces. This scheme enables static routes that are dependent on high-priority interfaces to be resolved and routing protocols to exchange information with peers over high-priority interfaces before the low-priority interfaces are replayed.

You can designate an IP or IPv6 interface as high priority either implicitly or explicitly:

- Implicit designation—Configure an IGP or PIM protocol on the interface.
- Explicit designation—Issue the **ip initial-sequence-preference 1** command on the IP subinterface, or the **ipv6 initial-sequence-preference 1** command on the IPv6 subinterface.

An IP or IPv6 interface can be designated as high priority by more than one protocol, the CLI command, or both. You can change an IP or IPv6 interface from high priority to low priority only by one of the following methods:

- Delete the IP or IPv6 interface.
- Remove all high-priority configuration from the IP or IPv6 interface, then reload the router.

ip initial-sequence-preference

ipv6 initial-sequence-preference

- Use to set the initial sequence preference value on an IP or IPv6 interface at the subinterface level.
- Specify a value of 1 to configure the interface as high-priority. Specify a value of 0 to configure the interface as low-priority.
- Example for IP


```
host1(config-subif)#ip initial-sequence-preference 1
```
- Example for IPv6


```
host1(config-subif)#ipv6 initial-sequence-preference 0
```
- There is no **no** version.
- See `ip initial-sequence-preference` and `ipv6 initial-sequence-preference`.

Upgrading Software

You cannot activate stateful SRP switchover when a different release of software is running on the standby SRP module. The router determines whether a release is the same by viewing the build date, the release filename, and the internal version number for the software on each SRP module.

The most efficient way to upgrade the software is to ensure that the standby SRP module is armed with the new release and then reload the standby SRP module. This reload occurs automatically after you download and arm a new release onto the active SRP module and the active SRP module subsequently synchronizes with the standby SRP module.

After reloading, and even though high availability mode is configured, the active SRP module reverts to using the file-system-synchronization operational mode for synchronizing updates. To complete the upgrade and place the system back in high-availability operational mode, you must execute the **srp switch** command to force the standby SRP module to take over as the active SRP module.



NOTE: Executing the **srp switch** command results in a cold restart of the router.

After the switchover is initiated, the formerly active SRP module reloads the software and starts running the same release as the newly active SRP module. When the formerly active SRP module becomes operational as the standby SRP module, the newly active SRP module detects that the release it is running is the same as that on the standby SRP module and allows the originally active SRP module to resume the high-availability operational mode.

If a fault occurs when the active SRP module is in file-system-synchronization operational mode, the standby SRP module detects the fault and takes over, and the router cold-restarts. For this reason, you must arm the new release *only* when you can accept the resulting window of vulnerability where high availability is disabled (that is, until the active and standby SRP modules are again running the same release).

Monitoring Stateful SRP Switchover

This section shows how to use the **show** commands to view your high availability configuration and how to clear the high availability switchover history for the router.

Stateful SRP Switchover show Commands

You can monitor various aspects of stateful SRP switchover using **show** commands. These aspects include redundancy modes and status, redundancy clients, historical information about redundancy on the router, and specific redundancy information for line modules and SRPs.

show ip interface
show ipv6 interface

- Use to display the initial sequence preference values as well as other information relating to the current state of all IP or IPv6 interfaces or the IP or IPv6 interfaces you specify.
- Field descriptions relating to redundancy. These fields are displayed when the **brief** keyword is not issued.
 - Operational initial sequence preference—Actual initial sequence preference as determined by the combination of the protocol configured on the interface

and the CLI configuration on the interface; the value is 0 (low priority) or 1 (high priority)

- Administrative—Configured initial sequence preference specified with the **ip initial-sequence-preference** command or the **ipv6 initial-sequence-preference** command; the value is 0 (low priority) or 1 (high priority)
- Example 1—Displays a summary of the redundancy status of an ERX14xx model

```
host1#show ip interface atm 2/0.1
```

```
ATM2/0.10 line protocol Atm1483 is up, ip is up
...
Warm-restart initial-sequence-preference: Operational = 1 Administrative
= 1
...
```

show redundancy

- Use to display the supported redundancy modes and other status relating to stateful SRP switchover. In particular, the output indicates any conditions that are preventing high availability from operating.
- Field descriptions
 - SRP
 - high-availability state—State of the high availability mode (disabled, active, or pending)
 - current redundancy mode—Redundancy mode currently being used by this router (high-availability or file-system-synchronization)
 - last activation type—Last type of activation that occurred on this router (that is, the method by which the SRP last booted [cold-switch or warm-switch])
 - Criteria Preventing High Availability from being Active—criteria required for high availability to be active.
 - Criteria Required for High Availability to be Active—criteria required for high availability to be active.



NOTE: All criteria for both options must be “yes” for high availability to be active.

- Line Card
 - automatic reverting—State of automatic reverting (on or off)
 - slots—Slots in which the line modules reside
 - hardware role—Function of the line module: primary or spare

- lockout config—Status of redundancy on this line module (protected—line module redundancy is enabled; locked out—line module redundancy is disabled)
 - backed up by slot—Slot that contains the line module that is a spare for this primary line module
 - sparing for slot—Slot that contains the primary line module for which this line module is a spare
 - revert at—Time at which you want line module to revert
 - midplane type—Identifier for the type of midplane
 - midplane rev—Hardware revision number of the redundancy midplane
 - fabric slice redundancy—Status of the fabric slice on the SRP modules or SFMs on the E120 and E320 routers
 - slot—Slot in which the fabric slice resides
 - slice state—State of the fabric slice (online, not present)
 - type—Identifier for the type of hardware (SRP module or SFM)
- Example 1—Displays a summary of the redundancy status of an ERX14xx model

```
host1#show redundancy
```

```
SRP
---
```

```
high-availability state: disabled
current redundancy mode: high-availability
last activation type: cold-switch
```

```
Criteria Preventing High Availability from being Active
```

criteria	met
Standby SRP is online and capable of mirroring?	No

```
Line Card
-----
```

```
automatic reverting is off
```

slot	hardware role	lockout config	backed up by slot	sparing for slot	revert at
3	---	---	---	---	---
8	spare	---	---	---	---
12	primary	protected	---	---	---

slots	midplane type	midplane rev
3	---	---
8	---	---
12	---	---

```

-----
8 - 13      6      0

```

- Example 2—Displays the redundancy status of an ERX14xx model in detail

```
host1#show redundancy detail
```

```
SRP
---
```

```

high-availability state: disabled
current redundancy mode: file-system-synchronization
last activation type:    cold-start

```

Criteria Required for High Availability to be Active

-----	-----	-----
	criteria	met
-----	-----	-----
Active SRP hardware supports High Availability?		Yes
High Availability mode configured?		No
Mirroring Subsystem present?		Yes
Mirroring activity levels within limits?		Yes
Network Core Dumps disabled?		Yes
Running configuration is safe for High Availability?		Yes
Standby SRP hardware supports High Availability?		Yes
Standby SRP is online and capable of mirroring?		Yes
Standby SRP is running the same release?		Yes

```
Line Card
-----
```

```
automatic reverting is off
```

slot	hardware role	lockout config	backed up by slot	sparing for slot	revert at
-----	-----	-----	-----	-----	-----
3	---	---	---	---	---
8	spare	---	---	---	---
12	primary	protected	---	---	---

slots	midplane type	midplane rev
-----	-----	-----
8 - 13	6	0

- See show redundancy.

show redundancy clients

- Use to display high availability clients and their various levels of high availability support.
- Specify an optional client type that you want to view (all, supported, unsafe, unsupported)



NOTE: Issuing this command without the optional client type results in showing only unsupported high availability clients (the default).

- Field descriptions
 - client—High availability client
 - mode—Whether the client is supported or unsupported for high availability
 - configuration—Safety level of the configuration based on whether or not the client is supported or unsupported and, in the case of those unsupported, whether or not the client has been configured. For example, if an unsupported client is configured on a router with high availability enabled, the configuration reads “unsafe”.
- Example 1

```
host1#show redundancy clients
```

Unsupported High Availability Clients

client	configuration
DHCP Proxy Client	safe
Global Ipv6	safe
IPsec Transport (ITM)	safe
l2tpDialoutGenerator	safe
DHCPv6 Local Server	safe
Radius Relay Server	safe

- Example 2

```
host1#show redundancy clients all
```

High Availability Client Information

client	mode	configuration
atm1483DataService	supported	safe
AA83	supported	safe
aaaServer	supported	safe
atmAal5	supported	safe
AAQS	supported	safe
atm	supported	safe
Bridged Ethernet	supported	safe
Transparent Bridging	supported	safe
dcm	supported	safe
dhcpExternal	supported	safe
DHCP Proxy Client	unsupported	safe
DS1	supported	safe
DS3	supported	safe
ethernet	supported	safe
Flow Inspection	supported	safe
frameRelay	supported	safe
FT1	supported	safe
Global Ipv6	unsupported	safe

Global Ip	supported	safe
HDLC	supported	safe
IKEP	supported	safe
ipflowstats	supported	safe
IpSubscriberManager	supported	safe
IPTU	supported	safe
IPVR	supported	safe
IPsec Transport (ITM)	unsupported	safe
l2tpDialoutGenerator	unsupported	safe
l2tp	supported	safe
LMGR	supported	safe
DHCPv4 Local Server	supported	safe
DHCPv6 Local Server	unsupported	safe
MPLS	supported	safe
PMGR	supported	safe
pppoe	supported	safe
ppp	supported	safe
qos	supported	safe
Radius Relay Server	unsupported	safe
RSVP	supported	safe
SCM	supported	safe
slothHelper	supported	safe
Cisco HDLC	supported	safe
ServiceManager	supported	safe
Sonet	supported	safe
SonetPath	supported	safe
SonetVT	supported	safe
IPsec Tunnel (ST)	supported	safe

- See show redundancy clients.

show redundancy history

- Use to display information about dates, times, and the number of occurrences for starts and switchovers.
- Use the **srp** keyword to view SRP module-specific information.
- Use the **detail** keyword to view additional history information.
- Field descriptions
 - system up time—Amount of time elapsed since the last cold boot
 - last cold start—Date and time the router experienced the last cold start
 - last cold switchover—Date and time the router experienced the last cold switchover
 - last warm switchover—Date and time the router experienced the last warm switchover
 - cold starts—Total number of cold starts the router has experienced
 - switchovers—Number of cold, warm, and consecutive warm switchovers the router has experienced
- Example 1

host1#show redundancy history

```

system up time:      0 00:08:01
last cold start:     2004-07-26 10:44:25
last cold switchover: 2004-07-25 18:51:56
last warm switchover: 2004-07-25 20:58:57

```

```

activation statistics:
  cold starts:      92
  switchovers:
    cold:           21
    warm:           147
  consecutive warm: 0

```

■ Example 2

host1#show redundancy history detail

```

system up time:      0 00:08:01
last cold start:     2004-07-26 10:44:25
last cold switchover: 2004-07-25 18:51:56
last warm switchover: 2004-07-25 20:58:57

```

```

activation statistics:
  cold starts:      92
  switchovers:
    cold:           21
    warm:           147
  consecutive warm: 0

```

SRP activation time	type	slot	system uptime	running release
-----	-----	----	-----	-----
2004-09-08 15:10:40	cold-start	00	---	erx_6-0-0b1-8.rel
2004-09-08 14:39:10	cold-start	00	---	erx_6-0-0b1-1.rel

■ See show redundancy history.

show redundancy line-card

- Use to display line-module-specific redundancy information.
- Field descriptions
 - automatic reverting—State of automatic reverting (on or off)
 - slots—Slots in which the line modules reside
 - hardware role—Function of the line module: primary or spare
 - lockout config—Status of redundancy on this line module (protected—Line module redundancy is enabled; locked out—Line module redundancy is disabled)
 - backed up by slot—Slot that contains the line module that is a spare for this primary line module

- sparing for slot—Slot that contains the primary line module for which this line module is a spare
- revert at—Time at which you want line module to revert
- midplane type—Identifier for the type of midplane
- midplane rev—Hardware revision number of the redundancy midplane
- Example

```
host1#show redundancy line-card
```

```
automatic reverting is off
```

slot	hardware role	lockout config	backed up by slot	sparing for slot	revert at
---	-----	-----	-----	-----	-----
3	---	---	---	---	---
8	spare	---	---	---	---
12	primary	protected	---	---	---

slots	midplane type	midplane rev
-----	-----	-----
8 - 13	6	0

- See show redundancy line-card.

show redundancy srp

- Use to display SRP-module-specific redundancy information.
- Use the brief keyword to display summary information about SRP redundancy.
- Field descriptions
 - SRP
 - high-availability state—State of the high availability mode (disabled, active, or pending)
 - current redundancy mode—Redundancy mode currently being used by this router (high-availability or file-system-synchronization)
 - last activation type—Last type of activation that occurred on this router (that is, the method by which the SRP last booted [cold start or warm start])
 - Criteria Required for High Availability to be Active—Criteria required for high availability to be active.



NOTE: All criteria must be “yes” for high availability to be active.

- Example 1


```

2004-07-25 18:58:01 warm-switch 06 0 00:12:01 L-07-25-60b1mrg-c.rel
2004-07-25 18:51:56 cold-switch 07 0 00:05:56 L-07-25-60b1mrg-c.rel
2004-07-25 18:46:54 cold-start 06 --- L-07-25-60b1mrg-c.rel
2004-07-25 17:44:48 warm-switch 06 0 00:14:32 L-07-25-60b1mrg-b.rel
2004-07-25 17:31:07 cold-start 07 --- L-07-25-60b1mrg-b.rel
2004-07-25 16:05:08 cold-start 07 --- L-07-25-60b1mrg-a.rel
2004-07-24 23:25:09 warm-switch 07 0 16:27:03 L-07-24-60b1mrg-b.rel
2004-07-24 23:18:23 cold-switch 06 0 16:20:17 L-07-24-60b1mrg-b.rel

```

- See `show redundancy switchover-history`.

Clearing the Redundancy History

You can use the **clear redundancy history** command to clear the stateful SRP switchover history for the router.

clear redundancy history

- Use to clear the detailed stateful SRP switchover history for the router.
- Example

```

host1# show redundancy history detail
system up time:      0 00:08:01
last cold start:     2004-07-26 10:44:25
last cold switchover: 2004-07-25 18:51:56
last warm switchover: 2004-07-25 20:58:57
activation statistics:
  cold starts:       92
  switchovers:
    cold:            21
    warm:            147
  consecutive warm:  0

```

SRP activation time	type	slot	system uptime	running release
2004-07-26 10:44:25	cold-start	07	---	L-07-25-60b1mrg-e.rel
2004-07-25 20:58:57	warm-switch	06	0 00:15:08	L-07-25-60b1mrg-e.rel
2004-07-25 20:53:41	warm-switch	07	0 00:09:51	L-07-25-60b1mrg-e.rel
2004-07-25 20:44:43	cold-start	06	---	L-07-25-60b1mrg-e.rel
2004-07-25 19:32:01	cold-start	06	---	L-07-25-60b1mrg-d.rel
2004-07-25 18:58:01	warm-switch	06	0 00:12:01	L-07-25-60b1mrg-c.rel
2004-07-25 18:51:56	cold-switch	07	0 00:05:56	L-07-25-60b1mrg-c.rel

```
host1# clear redundancy history
```

```

host1# show redundancy history
system up time:      0 00:08:01
last cold start:     2004-07-26 10:44:25
last cold switchover: 2004-07-25 18:51:56
last warm switchover: 2004-07-25 20:58:57
activation statistics:
  cold starts:       92
  switchovers:
    cold:            21
    warm:            147

```

```
consecutive warm: 0
SRP activation time      type      slot      system
-----             -----      ---      uptime      running release
```

- There is no **no** version.
- See clear redundancy history.

Chapter 4

Configuring a Unified In-Service Software Upgrade

This chapter describes how to prepare for and perform a unified in-service software upgrade (unified ISSU) of JUNOS Software on E120 and E320 Broadband Services Routers. A unified in-service software upgrade provides a way to upgrade to a higher-numbered release while minimizing the effect of the upgrade on traffic forwarded through the router.

- Unified ISSU Overview on page 55
- Unified ISSU Platform Considerations on page 57
- Unified ISSU Terms That Describe SRP and Line Module Behavior on page 57
- Unified ISSU References on page 58
- Unified ISSU Phases Overview on page 58
- Application Support for Unified ISSU on page 65
- Unexpected Application-Specific Behavior During Unified ISSU on page 73
- Before You Begin a Unified In-Service Software Upgrade on page 89
- Upgrading Router Software with Unified ISSU on page 91
- Halting the Unified ISSU Process and Restoring the Original State of the Router on page 95
- Monitoring a Unified In-Service Software Upgrade on page 97

Unified ISSU Overview

In software releases numbered lower than Release 6.0, all line modules are reloaded when an SRP switchover occurs. This reload disconnects user sessions and disrupts forwarding through the chassis. Stateful SRP switchover was introduced in JUNOS Release 6.0 to minimize the impact to the router of a stateful switchover from the active SRP module to the standby SRP module. Stateful SRP switchover (high availability) maintains user sessions during the switchover and data forwarding through the router continues to flow with little impact, thus improving the overall availability of the router.

The unified in-service software upgrade (unified ISSU) feature further extends router availability. Unified ISSU enables you to upgrade the router to a higher-numbered software release without disconnecting user sessions or disrupting forwarding through the chassis.

A conventional software upgrade—one that does not use the unified ISSU process—causes a router-wide outage for all users. Only static configurations (stored on the flash card) are maintained across the upgrade; all dynamic configurations are lost. A conventional upgrade takes 30-40 minutes to complete, with additional time required to bring all users back online.

When you perform a unified in-service software upgrade on a router that has one or more modules that do not support unified ISSU, these modules alone are upgraded by means of the legacy, conventional upgrade process. The unsupported modules undergo a cold reboot at the beginning of the unified ISSU process, and are held down until the in-service software upgrade is completed. Connections that pass through the unsupported modules are lost. The interfaces on these modules pass into a down state, which causes the physical layer and link layer to go down during the unified in-service software upgrade for those modules.

Applications that do not support unified ISSU applications cannot maintain state and configuration with minimal traffic loss across the upgrade to a higher-numbered release. When you attempt a unified in-service software upgrade on a router on which an ISSU-challenged application is configured, the unified in-service software upgrade cannot proceed. You must unconfigure the ISSU-challenged application to successfully perform the unified ISSU.

Router Behavior During a Unified In-Service Software Upgrade

The following behaviors are characteristic of a unified in-service software upgrade.

- Connections that were established before you begin the unified ISSU are maintained across the upgrade. Any such connection that was forwarding data continues to do so during and after the upgrade.
- New connections are denied until the upgrade is completed.
- Packet loss during the upgrade is limited. Bandwidth through the modules is reduced, but the impact is minimal.
- Graceful restart protocols do not time out during the unified ISSU.
- The unified in-service software upgrade has a minimal effect on the control and data planes. During the SRP module upgrade phase, forwarding through the fabric is interrupted for about 1 second on the E120 and E320 routers and about 4 seconds on the ERX1440 Broadband Services Router. During the line module upgrade phase, forwarding through the chassis is interrupted for about 15 seconds on the E120 and E320 routers and for about 50 seconds on the ERX1440 router.
- Diagnostic software is not run on any modules during a unified in-service software upgrade.
- The router undergoes a cold restart if you attempt to upgrade the software to a lower-numbered version with unified ISSU. The unified in-service software upgrade must be to a higher-numbered release than the running release.
- Additional memory is consumed during a unified in-service software upgrade. Available memory on a line module might not be sufficient due to the module's configuration. Unified ISSU can detect this limitation during the upgrade procedure and exit the process.

Unified ISSU Platform Considerations

Unified ISSU is supported on E120 and E320 routers. Unified ISSU is also supported on the ERX1440 router with the SRP-40G PLUS with 2GB of memory. Unified ISSU on the ERX1440 requires a license key.

For information about modules supported on E120 and E320 routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support unified ISSU.

For information about modules supported on the ERX1440 router:

- See *ERX Module Guide, Table 1, ERX Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support unified ISSU.

Redundant SRP modules are required for unified ISSU support.

Unified ISSU is not supported on the ERX7xx models, the ERX1410 router, and the ERX310 router.

Unified ISSU Terms That Describe SRP and Line Module Behavior

Table 6 on page 57 defines terms relevant to module behavior during a unified in-service software upgrade.

Table 6: Unified ISSU-Related Terms

Term	Meaning
Cold boot	The SRP module or line module loads diagnostics from the flash file system and runs them. When the diagnostics successfully complete, the operational image is loaded from the flash file system and then cold started.
Cold start	The SRP module or line module is initialized from the loaded operational image. The line modules are reloaded and the configuration is read from flash memory. When the line modules are operational, their configuration data is bulk downloaded and their interfaces become operational.
Cold restart	If the active SRP module fails, the standby SRP module takes the role of active SRP module. When high availability is not configured, the cold restart is similar to the cold start, except that the applications are already loaded into memory on the standby SRP module and ready to be started. The line modules are reloaded.

Table 6: Unified ISSU-Related Terms *(continued)*

Term	Meaning
Warm restart	If the active SRP module fails, the standby SRP module takes the role of active SRP module. Mirrored configuration data as well as any mirrored volatile data are already resident in memory. The line modules continue to forward data (with a small loss of packets when the fabric is switched from the formerly active SRP module to the newly active SRP module). The protocols and other applications re-initialize from the mirrored data and resynchronize communications with the line modules. When resynchronization is completed, the router resumes normal operations, including updates of any routing tables that result from changes that occurred during the warm restart.

Unified ISSU References

For more information about stateful SRP switchovers, see “Managing Stateful SRP Switchover” on page 25.

For more information about SRP module redundancy, see “Managing Module Redundancy” on page 7.

Unified ISSU Phases Overview

The JUNOS Software includes software modules that operate the following hardware components:

- SRP module
- Line module control plane
- Line module forwarding plane

A unified in-service software upgrade replaces the currently operating software on each of these components with a higher-numbered release. The unified ISSU also upgrades or re-creates the state and configuration data of the configured applications.

Before you begin the unified in-service software upgrade, you must first prepare the router for the upgrade. See “Before You Begin a Unified In-Service Software Upgrade” on page 89 for more information.

The unified in-service software upgrade takes place in three phases:

1. Initialization Phase—When you issue the **issu initialize** command, unified ISSU verifies whether all prerequisites for the upgrade have been met. The router is prepared for the upgrade. The configuration that has been mirrored to the standby SRP module is upgraded according to the upgrade release. This phase can last from a few minutes up to 40 minutes depending on the number of software releases across which the router is being upgraded.
2. Upgrade Phase—When you issue the **issu start** command, unified ISSU again verifies whether all prerequisites for the upgrade have been met. During this

phase the line module control plane and forwarding plane are upgraded and all three hardware components are resynchronized.

3. Service Restoration Phase—This phase automatically begins without user intervention when the upgrade phase has completed. During this final phase, the router is returned to a normal, runtime state.

The following sections describe these phases in more detail.

Unified ISSU Initialization Phase Overview

When you issue the **issu initialize** command, unified ISSU first verifies whether all requirements for the upgrade are met. The verification process examines the running release, the SRP modules, the line modules, line module redundancy, and the router configuration.

The **issu initialize** command does not interrupt or disrupt any of the runtime operations of the router. The command has no effect on changes of authorization, forwarding, or subscribers (except perhaps, rate of logins). You cannot manually change the file system redundancy mode from high availability to file synchronization until the unified in-service software upgrade is completed.



NOTE: We recommend that you make no configuration changes after you have issued the **issu initialization** command. As a best practice, ensure that your configuration is complete before you start the software upgrade.

During the initialization phase, you can halt the unified in-service software upgrade at any time and revert either to a normal SRP module switchover or to the previous state of the router. To stop the unified ISSU process, you must issue the **issu stop** command. If instead you simply exit the CLI session, the unified ISSU initialization phase continues.

The action taken when a requirement is not met depends on the requirement. For some failed verifications, the CLI warns you of the issue and prompts you to proceed or quit the upgrade process. More serious failures cause the CLI to exit the command and provide a message describing the issue.



NOTE: We recommend that you issue the **show issu** command before beginning the unified in-service software upgrade. The output of the command lists any necessary conditions that are not currently met on the router. You can therefore correct these failures before entering into the upgrade. You can issue the **show issu** command at any time.



NOTE: On E120 and E320 routers, you can hot swap an IOA during the initialization phase without affecting the in-service software upgrade. However, we strongly recommend that you perform any necessary hot swaps before you issue the **issu initialize** command.

If the standby SRP module reloads during the initialization phase, unified ISSU is halted. You must begin again by issuing the **issu initialize** command.

Application Data Upgrade on the Standby SRP Module

Synchronized modules can become unsynchronized because you can change the router configuration at any time until you issue the **issu start** command. When the verification process is completed, unified ISSU starts up the stateful SRP switchover process to maintain synchronization between the active SRP module and the standby SRP module during the SRP module upgrade phase.



NOTE: An SRP switchover from the active SRP module to the standby SRP module at this point in the unified in-service software upgrade causes a cold restart of the router because the SRP modules are running two different releases. The current release is on the active SRP module and the upgrade release is on the standby SRP module.

The application and configuration data that has been mirrored to the standby SRP module is upgraded as required by the new software release. The CLI displays the progress of the SRP module upgrade.

While data on the standby SRP module is upgraded, any new changes that are mirrored from the primary SRP module to the standby SRP module are also upgraded to the version required by the armed release.



NOTE: This process consumes significant CPU resources on the redundant SRP module and can take a considerable amount of time to complete. Performance of the active SRP module might be affected during the SRP module upgrade.

When the upgrade release has been synchronized to the standby SRP module, stateful SRP switchover is disabled until the unified in-service software upgrade is completed.

If you configure an application that does not support unified ISSU during the initialization phase, the initialization phase completes, but the **issu start** command subsequently fails.

SNMP Traps

When you issue the **issu initialization** command, the SNMP agent generates a `juniSystemIssuStateChange` trap with `juniSystemIssuState` set to `initializing`. When the unified ISSU initialization is completed, the SNMP agent generates a `juniSystemIssuStateChange` trap with `juniSystemIssuState` set to `initialized`.

Unified ISSU Upgrade Phase Overview

During the upgrade phase, the CLI supports only a reduced set of administrative commands. You cannot interrupt the upgrade phase. The upgrade phase cannot

commence if any CLI commands outside of this set are executing when you issue the **issu start** command.



NOTE: Although you can use any CLI session to issue the **issu start** command, we recommend that you issue the command from a session to the management console port. When the standby SRP module switchover takes place, all management network connections through the Ethernet management port are dropped, and you can access the router only through a console port until the service restoration phase is completed.

When you issue the **issu start** command, unified ISSU performs the following operations:

1. Verifies that the unified ISSU requirements on the router are still met.
2. Verifies that the active and standby SRP modules are synchronized. If they are not synchronized, forces a synchronization to ensure that all configuration and file system changes are propagated to the standby SRP module before proceeding with the upgrade.
3. Copies the NVS configuration from a backup memory area to the flash card on the standby SRP module. During the initialization phase, this configuration data was mirrored from NVS on the active SRP module and upgraded as required by the armed release.
4. Upgrades the control plane on each line module at the same time.
5. Switches control from the primary SRP module (running the current release) to the standby SRP module (running the armed upgrade release).
6. Upgrades the forwarding plane on each line module at the same time. The fabric is switched and upgraded.

You can view the status of the router and the progress of the upgrade at any time by issuing the **show issu** command from another terminal session to the router.



NOTE: While a unified ISSU operation is in progress, do not remove the SRP modules or attempt to reset them. Removing the SRP modules anytime during unified ISSU has an adverse impact.

After the unified ISSU operation is completed, issue the **show version** command. The output from a successful upgrade indicates the following:

- The formerly active SRP module has rebooted and come up as the new standby SRP module.
- The newly active SRP module indicates that the formerly active SRP has rebooted and has come up as standby SRP module

You can then remove an SRP module after issuing the **halt** command.

Exceptions During the Upgrade Phase

Table 7 on page 62 describes the behavior of the router during the upgrade phase when certain exceptional events take place outside the context of the unified in-service software upgrade.

Table 7: Router Response to Undesirable Events During the Upgrade Phase

Event	Router Action
The router reloads.	<ul style="list-style-type: none"> ■ The unified ISSU operation halts. ■ The router undergoes a cold restart. ■ The router boots with the armed upgrade release. ■ The line modules reboot.
The primary SRP module switches over to the standby SRP module.	<ul style="list-style-type: none"> ■ The unified ISSU operation halts. ■ The router undergoes a cold restart. ■ The router boots with the armed upgrade release. ■ The line modules reboot.
The standby SRP module reloads.	<ul style="list-style-type: none"> ■ The unified ISSU operation halts. ■ The router undergoes a cold restart. ■ The router boots with the armed upgrade release. ■ The line modules reboot.
A line module reloads.	<ul style="list-style-type: none"> ■ The line module is held down and prevented from rebooting until the service restoration phase is completed. The line module then undergoes a cold reboot to the running (pre-upgrade) release.
An IOA is hotswapped.	<ul style="list-style-type: none"> ■ Hot swapping is disabled during the upgrade phase. The line module undergoes a cold reboot and hot swapping is reenabled when the service restoration phase is completed.
An application that does not support unified ISSU is configured.	<ul style="list-style-type: none"> ■ This event can take place only before the issu start command is issued, because that command disables CLI configuration commands. When you issue the issu start command, after configuring such an application, the command exits and generates an error message.

Verification of Requirements

Because some time may have passed since unified ISSU verified the requirements for the upgrade during the initialization phase, unified ISSU re verifies all the same conditions.

Unified ISSU also verifies that no CLI configuration sessions are open, that no scripts or macros are running, and that any SNMP requests or CLI commands in progress complete within 5 seconds.

If any of the required conditions are not met, the CLI either exits the command with an error message or provides an informative message and prompts you to proceed or halt.

When all the conditions are met, the CLI prompts you to proceed. If you continue, then you can subsequently halt the upgrade only by reloading the router. If you exit the command, the router remains in the unified ISSU initialized state.

Upgrade Setup

At this stage all the preconditions have been met. The unified ISSU process shuts down all management interfaces to the router in order to prevent changes in the configuration.

Final preparation for the upgrade phase includes the following actions:

- **SNMP**—The SNMP agent generates a `juniSystemIssuStateChange` trap with `juniSystemIssuState` set to `upgrading` to indicate that the final phase of the operation has begun. When the trap is issued with this state value, the SNMP agent stops accepting any new SNMP gets or sets and does not issue any further traps.
- **CLI**—Most CLI commands are disabled. Only **reload**, **show issu**, and **show version** commands are available to you until the service restoration phase completes. These commands are available on the console and are not available to Telnet sessions.
- **External events**—Externally created events from sources other than the CLI are ignored. These events typically come from user connections, RADIUS servers, SRC software and SDX software, and SNMP, and include login requests, COA requests, multicast join requests, packet mirroring requests, and so on. Logout requests are cached and processed at a later stage.
- **Routing protocols**—The unified ISSU process prompts you to consider raising the link costs for each routing protocol that is configured on the router. Raising the link cost for routes through the upgrading router enables neighbors to recompute better routes to those destinations. If you choose to raise the link cost, the higher costs can take some time to propagate through the network. Because the router is unable to determine when this has completed, it waits for 2 minutes before proceeding to the next step in the upgrade.

The reason for raising the link cost is that when the upgrade of the line module control plane begins, routing protocol updates cannot be installed in the line modules until that upgrade completes. That period can be in the range 2–15 minutes. During the control plane upgrade, the routing protocols can still accept new routes and communicate with their neighbors but cannot install the routes.

- **Unsupported line modules**—Any unsupported line modules that are present are held down after the start of this phase when you can no longer gracefully exit from the unified ISSU process. The modules are held down for the duration of the unified in-service software upgrade and then undergo a cold boot to the original running release.
- **IGMP requests**—The router cannot handle IGMP requests for channel changing for IPTV implementations.

Line Module Arming

When the upgrade of the application data on the standby SRP upgrade is completed, unified ISSU temporarily arms the line modules with the upgrade release in a backup region of the memory.

Line Module Control Plane Upgrade

At this point, the upgrade release is preserved on each line module in some backup region. When signaled by the active SRP module, all supported line modules simultaneously reload and restart with the new release. Forwarding through the forwarding subsystem on the line modules—through the fabric of the system—is not affected by the reload.

The line modules then simultaneously recover any application data preserved in memory on the line module and upgrade that data into a format that is understood by the applications running on the new release. This operation can take in the range of 1–10 minutes depending on the size of the data and the upgrade path of the data. Each line module restores its operational state, running the new release with all data upgraded to a version acceptable to the new software.

If the upgrade process fails for any line module, that module undergoes a cold restart, but none of the other line modules is affected.

SRP Module Switchover

At this stage the primary SRP module is running the current release, the redundant SRP module is running the armed release, and the control plane on each supported line module is running the armed release.

When the primary SRP module has verified that all line modules are up, the redundant SRP module takes over control of the router by becoming the active SRP module. The primary, and formerly active, SRP module reboots to the armed release and serves as the standby SRP module.

All applications on the newly active SRP module are restarted. Each application reconstructs itself from the mirrored data, restoring its state and configuration as it was before the switchover. Forwarding through the fabric is interrupted for about 1 second on the E120 and E320 routers and about 4 seconds on the ERX1440 router.

The SRP switchover restarts the routing protocols and triggers a graceful restart because the routes need to be recomputed. This recalculation can take up to 90 seconds. Data continues to be forwarded through routes that were learned before the upgrade of the line module control planes.

Line Module Forwarding Plane Upgrade

While the applications on the SRP module and the line modules reconstruct themselves, they also begin to build up the new state for the forwarding subsystem. All applications on the line module signal the system when they are ready to start the forwarding upgrade.

Because upgrading the forwarding plane affects forwarding through the chassis for up to 30 seconds on the E120 and E320 routers and about 50 seconds on the ERX1440 router, unified ISSU does not proceed until the routing protocols have signaled that route reconvergence has completed. Unified ISSU then informs all line modules to simultaneously upgrade their forwarding subsystems.

The line modules then perform the following steps:

1. Halt forwarding through the line modules. Although any links to external devices remain up, incoming data is dropped.
2. Update any changed programmable hardware devices.
3. Update the forwarding subsystem with the new release and upgraded configuration data.
4. Update the routing tables with the reconverged routes.
5. Resume forwarding.

Unified ISSU Service Restoration Phase Overview

This is the final unified ISSU phase. At this point, all three major components of the router—the SRP modules, the line module control planes, and the line module forwarding planes—have been upgraded and forwarding has resumed through the chassis. The following actions take place during this phase:

- The CLI is re-enabled. All commands are made available to users.
- The SNMP agent is restarted and bulk statistics are collected and available for review. (The first interval of bulk statistics collection starts when unified ISSU is still in process. Therefore, the system performs bulk statistics collection after the first interval.)
- New login requests and logout requests are processed. The router begins to accept externally created events from sources other than the CLI and SNMP. These events typically come from user connections, RADIUS servers, and SRC software and SDX software, and include login requests, COA requests, multicast join requests, and so on.
- Logout requests that were cached at the start of the unified in-service software upgrade are processed.
- After the flash memory on the newly active SRP module is updated, stateful SRP switchover is available to the router.

At this point the unified in-service software upgrade is completed, and the router is restored to normal operation. Any line modules that reloaded during the upgrade phase and were therefore held down are now rebooted to the original running release.

Application Support for Unified ISSU

When an application supports unified ISSU, you can configure the application on the router and proceed with the unified in-service software upgrade with no adverse impact to the upgrade.

Applications that do not support unified ISSU cannot maintain state and configuration with minimal traffic loss across the upgrade. When you attempt the unified in-service software upgrade on a router that is configured with an ISSU-challenged application, the unified in-service software upgrade is halted and cannot proceed unless you remove the configuration. An application that does not support high availability cannot support unified ISSU.

Table 8 on page 66 indicates which applications support or do not support a unified in-service software upgrade, as well as limitations on their behavior.

Table 8: Application Support for Unified In-Service Software Upgrades

Application	Supported	Unsupported	Notes
Physical Layer Protocols			
DS1 (E120 and E320)	✓	–	–
DS1 (ERX1440)	–	–	–
DS3	✓	–	–
HDL	✓	–	–
SONET/SDH	✓	–	<p>Unified ISSU support is provided only for non-channelized APS IOAs. Also, unified ISSU can proceed only if you have not configured APS on the OCx/STMx ATM or OCx/STMx POS line modules. If you have configured APS, a warning message is displayed and the router cannot proceed with the unified ISSU.</p> <p>The unified ISSU process for channelized line modules remains unchanged.</p> <p>E120 and E320 routers do not support APS.</p>
SONET/SDH VT	–	✓	–
Link-Layer Protocols			

Table 8: Application Support for Unified In-Service Software Upgrades *(continued)*

Application	Supported	Unsupported	Notes
ARP	✓	–	ARP entries in the ARP cache do not time out because no ARP aging occurs during unified ISSU. When the unified ISSU is completed, the ARP cache contains the same entries as it had before the unified ISSU began.
ATM	✓	–	–
ATM 1483 bulk configuration of dynamic interfaces	✓	–	–
ATM bulk configuration of static interfaces	✓	–	–
Bridged Ethernet	✓	–	–
Cisco HDLC	✓	–	–
Ethernet (with and without VLANs)	✓	–	–
Frame Relay	–	–	–
PPP	✓	–	–
PPPoE	✓	–	–
Transparent bridging	✓	–	–
Unicast Routing			
Access Routes	✓	–	–
BGP	✓	–	–
FTP	✓	–	Although unified ISSU supports FTP in active state, no file transfer operation can be in progress while performing unified ISSU.
IP	✓	–	–
IPv6	–	✓	Unified ISSU does not support IPv6.

Table 8: Application Support for Unified In-Service Software Upgrades *(continued)*

Application	Supported	Unsupported	Notes
IPSec Transport (E120 and E320)	Not applicable	Not applicable	E120 and E320 routers do not support IPSec.
IPSec Transport (ERX1440)	–	–	–
IPSec Tunnels (E120 and E320)	Not applicable	Not applicable	E120 and E320 routers do not support IPSec.
IPSec Tunnels (ERX1440)	–	–	–
IS-IS	✓	–	Support only when graceful restart is configured.
OSPF	✓	–	Support only when graceful restart is configured.
RIP	✓	–	–
Static Routes	✓	–	–
Telnet	✓	–	Authentication and command authorizations on Telnet sessions fail during the upgrade phase and Telnet sessions are dropped.
IPv4 Multicast Routing			
Multicast Routing	✓	–	–
ANCP (L2C)	✓	–	Unified ISSU can proceed if ANCP is configured. However, ANCP has no graceful restart extensions and therefore it cannot maintain its dynamic state across the upgrade. Consequently, all ANCP sessions are brought down and then restored when the upgrade is completed.

Table 8: Application Support for Unified In-Service Software Upgrades (continued)

Application	Supported	Unsupported	Notes
DVMRP (E120 and E320)	✓	–	–
DVMRP (ERX1440)	–	–	–
IGMP	✓	–	–
PIM	✓	–	–
IPv6 Multicast Routing			
Multicast Routing	–	✓	Unified ISSU does not support IPv6.
MLD	–	✓	Unified ISSU does not support IPv6.
PIM	–	✓	Unified ISSU does not support IPv6.
Multiprotocol Label Switching			
MPLS	✓	–	–
BGP signaling	✓	–	–
LDP signaling	✓	–	–
RSVP-TE signaling	✓	–	–
Local cross-connects between layer 2 interfaces using MPLS	✓	–	–
Policies and QoS			
Policies	✓	–	–
QoS	✓	–	–
Remote Access			
AAA	✓	–	The following configuration is not supported: The subscriber username and password are on an ATM circuit in Bridged Ethernet over ATM or IP over ATM configurations.
DHCP External Server and Packet Trigger	✓	–	–

Table 8: Application Support for Unified In-Service Software Upgrades *(continued)*

Application	Supported	Unsupported	Notes
DHCP Packet Capture	✓	–	Configuration of DHCP packet capture does not prevent unified ISSU from proceeding. However, the capturing of packets on the line modules is halted when the unified ISSU upgrade phase commences. Packet capture resumes automatically during the unified ISSU service restoration phase.
DHCP Proxy Client	–	✓	–
DHCP Relay Proxy	–	✓	DHCP relay proxy continues processing of DHCP release requests during the unified ISSU to maintain server-client synchronization. State is preserved across the upgrade; statistics are not preserved.
DHCP Relay Server	✓	–	–
DHCPv4 Local Server	✓	–	Forwarding outages that take place during a unified ISSU can affect DHCP lease renewal. Before you begin unified ISSU, we recommend that you configure the DHCP local server address pools with a minimum lease time of 120 minutes to ensure that leases do not expire during the upgrade.
DHCPv6 Local Server	–	✓	Unified ISSU does not support IPv6.

Table 8: Application Support for Unified In-Service Software Upgrades *(continued)*

Application	Supported	Unsupported	Notes
L2TP	✓	–	Unified ISSU forces an L2TP failover for all established tunnels. L2TP failover resynchronization is required for successful recovery of a tunnel and its sessions following the upgrade.
L2TP Dialout	–	✓	–
Local Address Pools	✓	–	–
Local Authentication Server	✓	–	–
RADIUS Client	✓	–	–
RADIUS Dynamic-Request Server	✓	–	–
RADIUS Initiated Disconnect	✓	–	–
RADIUS Relay Server	–	✓	–
RADIUS Route-Download Server	✓	–	–
SRC Client	✓	–	–
Service Manager	✓	–	–
Subscriber Manager	✓	–	–
TACACS +	✓	–	–
Miscellaneous			
Bulk statistics	✓	–	–
Denial of Service (DoS) protection	✓	–	–
HTTP server	✓	–	–
IOA hot swap	–	✓	–
J-Flow (IP flow statistics)	✓	–	–

Table 8: Application Support for Unified In-Service Software Upgrades *(continued)*

Application	Supported	Unsupported	Notes
Line Module Redundancy	✓	–	<p>You can use the active spare line module for unified ISSU operations. You do not have to revert to the primary line module. The following sets of line modules and IOAs are supported:</p> <ul style="list-style-type: none"> ■ ATM: OC3-4A, OC3/OC12/DS3-ATM ■ POS: OC3-4P ■ Line Modules <ul style="list-style-type: none"> ■ ES2 4G LM ■ ES2 10G Uplink LM ■ ES2 10G LM ■ ES2 10G ADV LM ■ IOAs <ul style="list-style-type: none"> ■ ES2-S1 GE-4 IOA ■ ES2-S1 GE-8 IOA ■ ES2-S3 GE-20 IOA ■ ES2-S1 10GE IOA ■ ES2-S2 10GE PR IOA ■ ES2-S1 OC3-8 STM1 ATM IOA ■ ES2-S1 OC12-2 STM4 ATM IOA ■ ES2-S1 OC12-2 STM4 POS IOA ■ ES2-S1 OC48 STM16 POS IOA
Mobile IP Home Agent	–	✓	–

Table 8: Application Support for Unified In-Service Software Upgrades *(continued)*

Application	Supported	Unsupported	Notes
Network Address Translation (NAT)	✓	–	You must remove the NAT license configuration as well as the NAT configuration from the router.
NTP	✓	–	–
Resource Threshold Monitor	✓	–	–
Response Time Reporter	✓	–	–
Route Policy	✓	–	–
SNMP	✓	–	–
Subscriber Interfaces	✓	–	–
Tunnels (GRE and DVMRP)	✓	–	–
VRRP	✓	–	–

Unexpected Application-Specific Behavior During Unified ISSU

This section describes the behavior of applications that vary from the expected behavior during a unified in-service software upgrade.

- “AAA Authentication and Authorization Disabled” on page 74
- “ATM Affected Behaviors” on page 74
- “DHCP Affected Behaviors” on page 75
- “DoS Protection State Freeze” on page 75
- “Ethernet Affected Behaviors” on page 76
- “FTP Server File Transfer Behaviors” on page 77
- “IS-IS Effects on Graceful Restart and Network Stability” on page 79
- “L2TP Failover of Established Tunnels” on page 81
- “OSPF Effects on Graceful Restart, Timeouts, and Network Stability” on page 82
- “PIM Suspended During Unified ISSU” on page 84
- “Subscriber Logins and Logouts Suspended During Unified ISSU” on page 84
- “SONET/SDH Behavior During Unified ISSU” on page 85

- “T3” on page 85
- “TACACS+ Services Not Available” on page 85
- “Interruption in Traffic Forwarding for Layer 3 Routing and Signaling Protocols” on page 86
- “Recommended Routing Protocol Timer Settings” on page 88

AAA Authentication and Authorization Disabled

Authentication and command authorization are temporarily disabled on the serial console connection during the upgrade phase. You can connect to the serial console and issue the **reload**, **show issu**, and **show version** commands without the action of authentication and authorization servers, such as RADIUS or TACACS+.

ATM Affected Behaviors

The following aspects of ATM behavior during unified ISSU are different than the behavior during normal router operations.

ILMI Sessions Not Maintained

The router does not maintain ILMI sessions during a unified in-service software upgrade. The router terminates all ILMI sessions and restarts them during the upgrade. If the ILMI protocol is enabled on any port, you are warned during the initialization phase when unified ISSU is verifying the prerequisites for the upgrade. You can choose to continue the upgrade—and bring down the sessions—or to halt the unified in-service software upgrade.

OAM CC Effects on VCC

When an ATM VC is configured as an OAM CC source, periodic OAM cells are generated for about 15 seconds. The device configured as the OAM CC sink is likely to declare the VCC down during this time. Unified ISSU generates a warning when it detects an OAM CC source configuration during the initialization phase while unified ISSU is verifying the prerequisites for the upgrade. You can choose to continue or halt the unified in-service software upgrade.

When an ATM VC is configured as OAM CC sink, it cannot receive OAM CC cells generated by the source for about 15 seconds. The OAM CC does not time out and the VCC is not declared down. OAM CC cell generation resumes when the unified ISSU operation is completed.

OAM VC Integrity Verification Cessation

During the unified ISSU operation, verification of OAM VC integrity stops. This verification resumes when the unified ISSU operation is completed.

ATM does not respond to incoming OAM loopback cells during the upgrade for a period of less than 30 seconds. The lack of response might cause ATM peers to declare VCC (VPC) down.

Port Data Rate Monitoring Cessation

The monitoring of ATM port data rates is halted during a unified in-service software upgrade. Monitoring resumes immediately after the unified ISSU operation is completed. The data rates reported by the **show atm interface** command are inaccurate for the period of one configured load interval after unified ISSU is completed.

VC and VP Statistics Monitoring Halts Unified ISSU Progress

A unified in-service software upgrade cannot proceed if VC or VP statistics monitoring is in progress.

DHCP Affected Behaviors

DHCP Common Component Information Suspended

The common DHCP component supports unified ISSU. This component configures option 60 vendor-option strings when you issue the **set dhcp vendor-option** command. The DHCP common component ceases all CLI and SNMP operations when you issue **issu start** command. You can therefore obtain no information about the common DHCP infrastructure until the unified in-service software upgrade is completed.

DHCP Relay and DHCP Relay Proxy Prevent Unified ISSU

The DHCP relay and DHCP relay proxy applications do not support unified ISSU. You must completely unconfigure these applications from all virtual routers to perform a unified in-service software upgrade.

DHCP Packet Capture Halted on Line Modules

The DHCP packet capture application supports unified ISSU in that its configuration does not halt a unified in-service software upgrade. However, packet capture on line modules is halted during the upgrade phase. Packets are not captured and buffered for later forwarding to the SRP module during this phase. Capture resumes automatically during the service restoration phase.

DoS Protection State Freeze

The denial-of-service (DoS) protection application freezes its state when the in-service software upgrade is initiated. Any suspicious control flow, protocol, or priority remains suspicious until unified ISSU completes.

Freezing the DoS protection state prevents any active control flows from interfering with the system while the unified ISSU is in progress. However, no new control flows, protocols, or priorities are monitored for suspicious activity, and no suspicious activity can be detected until the upgrade is completed.



NOTE: Because of this limitation on DoS functionality, we recommend that you do not initiate unified ISSU until all suspicious control flows, protocols, and priorities have been resolved.

When the unified in-service software upgrade is completed, the DoS protection application resumes attending to all dynamic state that was frozen at the beginning of the unified ISSU process.

Some suspicious control flows might remain in a suspicious list based on your configuration, if the upgrade software version has DoS protection classification algorithms that are better or different than in the previous version. Because flows are discovered and monitored at 1-second intervals, the new conditions might cause these flows to be removed. You do not need to explicitly clear the flows when unified ISSU is completed.

Ethernet Affected Behaviors

The following aspects of Ethernet behavior during a unified in-service software upgrade are different than during normal router operations.

ARP Packets Briefly Not Sent or Received

There is a brief period at the end of the upgrade phase when ARP packets are not sent or received. This event can affect ARP entries on attached devices that were in the process of being aged out.

Link Aggregation interruption

During the unified in-service software upgrade, LACP PDUs are not generated or received for about 15 seconds on Ethernet ports configured with LACP.

This interruption has no effect on the local side of the link because JUNOS Software generates LAC PDU packets every 30 seconds. The link is not declared down unless packets are missed three times. LACP packet generation continues when the unified ISSU operation is completed.

If a device on the other end of the link is configured with the short timeout of one second, then the device is likely to declare the link to be down and remove the link from the LAG bundle.

Port Data Rate Monitoring Halted

The monitoring of Ethernet port data rate is halted during a unified in-service software upgrade. Monitoring resumes immediately after the unified ISSU operation is completed. The data rates reported by the **show interface** command are inaccurate for the period of one configured load interval after unified ISSU is completed.

VLAN Statistics Monitoring Halts Unified ISSU Progress

A unified in-service software upgrade cannot proceed if VLAN statistics monitoring is in progress.

FTP Server File Transfer Behaviors

You can perform the unified ISSU operation even when the FTP server is enabled on the router. However, no file transfer process, such as uploading or downloading of files, creating of directories, or removing of files, can be in progress to enable the unified ISSU operation to complete successfully.

When you issue the **issu initialize** command, unified ISSU checks for open FTP connections or active file transfer sessions. At this stage, existing connections are not terminated and new connections can also be established. When you issue the **issu start** command, all FTP connections, including data and control connections, are disconnected. Although the listening port is still available at this stage, any attempt to create a new connection and incomplete file operations on existing connections fail with an appropriate error message from the FTP server.

The **issu start** command is not executed if file transfer operations are in progress. You must issue the **ftp-server flush** command to forcibly terminate the file transfer process. When you are prompted to confirm, type **y** to confirm to close all active file transfer jobs.



CAUTION: Because using the **ftp-server flush** command causes a forced and ungraceful termination of all file transfer jobs that are in progress to start the unified ISSU process, use it only when it is absolutely necessary. We recommend that you wait for file transfer operations that are in progress to complete gracefully before you perform unified ISSU, if your situation enables you to do so.

The following example shows the output of the **show ftp-server** command in a scenario where FTP server is enabled, but no open file transfer connections exist.

```
host1#show ftp-server
```

```
FTP Server state: enabled, 0 open connections
Statistics since server was last started:
  attempts: 3
  failed hosts: 0
  failed users: 0
Statistics since last system reload:
  attempts: 3
  failed hosts: 0
  failed users: 0
```

To display detailed information about unified ISSU status and warnings in addition to criteria required for unified ISSU and whether the router hardware and software meet the required criteria, issue the **show issu detail** command.

```
host1#show issu detail
```

```
ISSU state:      idle
ISSU description: ISSU is currently idle
criteria met:    No, upgrade error(s) found
running release: dtnguyen.rel
armed release:   dtnguyen.rel
```

#	ISSU Criteria Summary	Met
1	In-Service Software Upgrade ready?	Yes
2	High-Availability ready?	No
3	Line modules ready?	Conditional
4	Configuration conversion support ready?	Yes
5	CLI sessions ready?	Yes
6	Routing applications ready?	Yes
7	Protocol timers ready?	Yes

The following example shows a case when a few clients are connected to the FTP server, and the FTP ISSU state becomes conditional. However, unified ISSU begins without any error. All existing connections are dropped when you issue the **issu start** command and the upgrade runs.

```
host1#show ftp-server
```

```
FTP Server state: enabled, 1 open connections
Statistics since server was last started:
  attempts: 3
  failed hosts: 0
  failed users: 0
Statistics since last system reload:
  attempts: 3
  failed hosts: 0
  failed users: 0
```

```
host1#show issu detail
```

```
ISSU state:      idle
ISSU description: ISSU is currently idle
criteria met:    No, upgrade error(s) found
running release: dtnguyen.rel
armed release:   dtnguyen.rel
```

#	ISSU Criteria Summary	Met
1	In-Service Software Upgrade ready?	Yes
2	High-Availability ready?	No
3	Line modules ready?	Conditional
4	Configuration conversion support ready?	Yes
5	CLI sessions ready?	Yes
6	Routing applications ready?	Conditional
->	Criteria: There are open FTP connections Impact: Open connections will be disconnected during ISSU process Remedy: Close all FTP sessions Reporting slot: 7	Conditional
7	Protocol timers ready?	Yes

The following example shows when an ongoing file transfer operation is detected during the initialization phase or validation phase. In this case, the prerequisite verification that unified ISSU performs fails. Unified ISSU does not proceed until the active file transfer operations are terminated. Issue the **ftp-server flush** command to forcibly terminate all FTP sessions.

```
host1#show ftp-server
FTP Server state: enabled, 1 open connections
Statistics since server was last started:
    attempts: 3
    failed hosts: 0
    failed users: 0
Statistics since last system reload:
    attempts: 3
    failed hosts: 0
    failed users: 0
```

```
host1#show issu detail
ISSU state: idle
ISSU description: ISSU is currently idle
criteria met: No, upgrade error(s) found
running release: dtnguyen.rel
armed release: dtnguyen.rel
#          ISSU Criteria Summary          Met
--  -----
1  In-Service Software Upgrade ready?      Yes
2  High-Availability ready?                No
3  Line modules ready?                     Conditional
4  Configuration conversion support ready?  Yes
5  CLI sessions ready?                     Yes
6  Routing applications ready?              No
-> Criteria: FTP file transfer is in progress  No
    Impact: ISSU cannot be performed when file transfer is in pr
    ogress
    Remedy: Abort transfer with "ftp-server flush" or wait until
    transfer is done
    Reporting slot: 7
7  Protocol timers ready?                  Yes
```

```
host1#ftp-server flush
This command will terminate all FTP sessions, continue? [confirm]
host1#
```

New FTP connections are not allowed and all existing FTP connections are dropped after the unified ISSU process begins. Also, no remote file operations are allowed while unified ISSU is in progress. If unified ISSU is aborted, FTP server is returned to the state in which it was before unified ISSU was started.

IS-IS Effects on Graceful Restart and Network Stability

IS-IS has the following issues to consider before you begin a unified in-service software upgrade:

- Graceful restart—Required
- Routing around the upgrading router—Optional

Configuring Graceful Restart Before Unified ISSU Begins

You must configure IS-IS graceful restart on the router and on all IS-IS neighbors before you begin the unified in-service software upgrade. When the unified ISSU process verifies the upgrade requirements during the initialization phase, it detects whether graceful restart is configured. If it is not configured, the CLI displays a warning message and prompts you to proceed or halt. You can stop at this point to configure graceful restart.

If instead you proceed, the unified in-service software upgrade can complete successfully, but the IS-IS neighbors are likely to break the adjacencies with the upgrading router and consider that routes formerly reached through this router are now unreachable. When the unified in-service software upgrade completes and the routing protocols restart, the IS-IS neighbors can relearn the routes through the router.

When you issue the **issu start** command, IS-IS lengthens its hello timer values and sends LSPs with the new values. The upgrade proceeds when the IS-IS neighbors have acknowledged the new values.

Configuring Graceful Restart When BGP And LDP Are Configured

When BGP, IS-IS, and LDP are all configured on a router on which you want to perform a unified in-service software upgrade, ensure that the IS-IS graceful restart timeout is longer than the LDP graceful restart timeout. The IS-IS graceful restart does not complete when the LDP graceful restart timeout is longer than the IS-IS graceful restart timeout. Configure IS-IS graceful restart timeout with the **nsf t3** command. Configure LDP graceful restart timeout with the **mpls ldp graceful-restart timers max-recovery** command.

Routing Around the Restarting Router to Minimize Network Instability



NOTE: The situation described in this section is very uncommon. This rare circumstance arises when you have redundant uplinks to the core and network topology changes cause routes to go through the upgrading router. In a typical network design, this is not an issue and you do not need to route peers around the upgrading router,

During the unified ISSU upgrade phase, network instability can result if the restarting router goes into an unstable state after the unified ISSU process fails. Some IS-IS traffic loss occurs during the resulting line module resets. For those reasons, you might want IS-IS peers to route around the router that is being upgraded.

You can use the **overload advertise-high-metric issu** command to cause the router to advertise a high metric to its neighbors so that they route around the upgrading router. When you issue the **issu start** command, the router raises the metric to the maximum link cost on all interfaces running IS-IS. The maximum value depends on

the metric type. IS-IS neighbors then choose a path with lower metrics to reach any destination that was previously reached through the upgrading router. When unified ISSU is completed, IS-IS reverts the metrics back to the values that were configured before the unified in-service software upgrade.

When traffic engineering has been configured, the traffic engineering metrics are also increased. New tunnels are not established through the upgrading router and any tunnels undergoing re-optimization in other routers go around the upgrading router.

IS-IS support for unified ISSU does not depend on this configuration. If you do not issue the **overload advertise-high-metric issu** command, the unified in-service software upgrade can still proceed to successful completion without disrupting IS-IS functionality.

overload advertise-high-metric issu

- Use to cause IS-IS to advertise the maximum link metric on all interfaces to IS-IS neighbors when a unified in-service software upgrade is started.

- Example

```
host1(config-router)#overload advertise-high-metric issu
```

- Use the **no** version to send the configured link costs to neighbors during the unified in-service software upgrade.
- See `overload advertise-high-metric issu`.

L2TP Failover of Established Tunnels

L2TP never declares itself as unified ISSU unsafe. However, unified ISSU forces an L2TP failover for all established tunnels. Successful recovery of a tunnel and its sessions following the unified in-service software upgrade requires a successful L2TP failover resynchronization, either by the L2TP silent failover method or the L2TP failover protocol.

When the L2TP silent failover method is configured on ERX1440 router, use the **l2tp retransmission** command to set the retransmission retry count to 8 for the remote peers. A value of more than 7 helps ensure that the remote peers keep retransmitting control messages for the duration of the unified ISSU warm restart and the tunnels are not disconnected.

See *Specifying the Number of Retransmission Attempts*.

When the unified ISSU operation attempts to verify the upgrade prerequisites, a warning message is generated if any tunnels are present for which failover resynchronization is disabled.

You can use the **show l2tp tunnel failover-resync disable** command to identify the tunnels referred to by the warning message. The command enables filtering based upon the effective failover resynchronization mechanism:

```
host1#show l2tp tunnel failover-resync disable
```

```
L2TP tunnel 2/1 is Up with 1 active session
1 L2TP tunnel found
```

If a successful failover resynchronization cannot be performed for a tunnel following the upgrade, then the tunnel and all of its sessions are subject to disconnection.

L2TP automatically detects a peer L2TP disconnect after the unified in-service software upgrade is completed by detecting a control channel failure.

When peer LNSs are not configured with PPP keepalives or inactivity timeouts, you must configure an inactivity timeout for L2TP on the LAC. This timeout enables the router to detect a PPP disconnect when signaling has been dropped during the unified ISSU forwarding interruption. In the absence of this configuration, the connection at the LAC and LNS is left as logged in for an extended period of time following the upgrade.

OSPF Effects on Graceful Restart, Timeouts, and Network Stability

OSPF has the following issues to consider before you begin a unified in-service software upgrade:

- Graceful restart—Required
- Dead interval—Required
- Routing around the upgrading router—Optional

Configuring Graceful Restart Before Unified ISSU Begins

You must configure OSPF graceful restart before you begin the unified in-service software upgrade. When the unified ISSU process verifies the upgrade requirements during the initialization phase, it detects whether graceful restart is configured. If it is not configured, the CLI displays a warning message and prompts you to proceed or halt. You can stop at this point to configure graceful restart.

If instead you proceed, the unified in-service software upgrade can complete successfully, but the OSPF neighbors are likely to break the adjacencies with the upgrading router and consider that routes formerly reached through this router are now unreachable. When the unified in-service software upgrade completes and the routing protocols restart, the IS-IS neighbors can relearn the routes through the router.

You must also ensure that the OSPF neighbors have been configured as graceful restart helper routers. During the unified ISSU initialization phase, OSPF graceful restart on the upgrading router cannot verify whether its neighbors are helper routers, and reports that fact by means of the CLI.

Configuring Graceful Restart When BGP And LDP Are Configured

When BGP, LDP, and OSPF are all configured on a router on which you want to perform a unified in-service software upgrade, ensure that the OSPF graceful restart timeout is longer than the LDP graceful restart timeout. The OSPF graceful restart does not complete when the LDP graceful restart timeout is longer than the OSPF graceful restart timeout. Configure OSPF graceful restart timeout with the

graceful-restart restart-time command. Configure LDP graceful restart timeout with the **mpls ldp graceful-restart timers max-recovery** command.

Configuring a Longer Dead Interval Than Normal

To prevent OSPF from timing out to the OSPF neighbors, you must configure a dead interval that is longer than the expected forwarding outage for the platform. During the initialization phase, unified ISSU displays the recommended dead interval in a warning message. For information about the expected forwarding outage, see “Interruption in Traffic Forwarding for Layer 3 Routing and Signaling Protocols” on page 86

Routing Around the Restarting Router to Minimize Network Instability



NOTE: The situation described in this section is very uncommon. This rare circumstance arises when you have redundant uplinks to the core and network topology changes cause routes to go through the upgrading router. In a typical network design, this is not an issue and you do not need to route peers around the upgrading router,

During the unified ISSU upgrade phase, network instability can result if the restarting router goes into an unstable state after the unified ISSU process fails. Some OSPF traffic loss occurs during the resulting line module resets. For those reasons, you might want OSPF peers to route around the router that is being upgraded.

You can use the **overload advertise-high-metric issu** command to cause the router to advertise a high link cost to its neighbors so that they route around the upgrading router. When you issue the **issu start** command, the router raises the link cost to the maximum link cost on all interfaces running OSPF. The higher cost is advertised in the OSPF LSAs. OSPF neighbors then choose a path with lower metrics to reach any destination that was previously reached through the upgrading router. When unified ISSU is completed, OSPF reverts the link costs back to the values that were configured before the unified in-service software upgrade.

When traffic engineering has been configured, the traffic engineering metrics are also increased. New tunnels are not established through the upgrading router and any tunnels undergoing re-optimization in other routers go around the upgrading router.

OSPF support for unified ISSU does not depend on this configuration. If you do not issue the **overload advertise-high-metric issu** command, the unified in-service software upgrade can still proceed to successful completion without disrupting OSPF functionality.

overload advertise-high-metric issu

- Use to cause OSPF to advertise the maximum link cost on all interfaces to OSPF neighbors when a unified in-service software upgrade is started.
- Example

```
host1(config-router)#overload advertise-high-metric issu
```

- Use the **no** version to send the configured link costs to neighbors during the unified in-service software upgrade.
- See overload advertise-high-metric issu.

PIM Suspended During Unified ISSU

You can minimize PIM traffic loss during the unified in-service software upgrade by issuing the **ip pim dr-priority** command to set a priority so that PIM neighbors do not forward traffic through the upgrading router. By default, a PIM interface has a priority of one. If you set the priority to one, the lowest possible priority, then the upgrading router is not selected to be a designated router in the PIM network if an interface on another router in that network has a higher priority.

ip pim dr-priority

- Use to set a priority by which a router is likely to be selected as the designated router.
- Example

```
host1(config-if)#ip pim dr-priority 1
```

- Use the **no** version to restore the default value, 1.
- See ip pim dr-priority.

Subscriber Logins and Logouts Suspended During Unified ISSU

All new subscriber logins are ignored during the upgrade phase. New subscriber logouts are cached and processed after the unified ISSU operation is completed.

Subscriber Statistics Accumulation or Deletion

All subscriber statistics present in the line modules are cleared when the line module forwarding planes are upgraded. For this reason, the router has to read the statistics from the forwarding plane before it is upgraded.

However, forwarding through the line modules continues after that point, until the line module forwarding plane is upgraded. Some statistics can therefore accumulate in the forwarding plane in the interval between these two events. These statistics are not preserved across the upgrade.

Statistics for subscribers that log out during the forwarding plane upgrade are collected and reported before the forwarding plane is reloaded. Statistics are not collected for any subscribers who are connected before you issue the **issu start** command but who log out before the forwarding plane upgrade is completed.

The following subscriber statistics are preserved across the upgrade:

- All policy statistics
- Accounting statistics reported by IP: in bytes, out bytes, in packets, out packets

- Accounting statistics reported by L2TP: in octets, out octets, in packets, out packets
- Accounting statistics reported by PPP: in octets, out octets, in packets, out packets

All other statistics are set to zero, including all statistics belonging to the SNMP generic interface MIB for every interface layer.

SONET/SDH Behavior During Unified ISSU

During a unified in-service software upgrade, several aspects of SONET behavior differ from normal operation.

- SONET APS is supported only for non-channelized APS IOAs.

During a unified in-service software upgrade, if you have configured APS functionality on the non-channelized APS IOAs, the unified ISSU process fails and a warning message is displayed. If you have not configured APS functionality, the unified ISSU process succeeds and the line modules (OC3/OC12) do not get rebooted.



NOTE: The unified ISSU process for the channelized APS IOAs has not been modified. The channelized APS IOAs are rebooted during a unified in-service software upgrade.

- During a conventional software upgrade, a SONET loss-of-signal defect lasts more than 2.5 seconds, causing the router to declare an LOS failure. Devices on the remote end of SONET links detect the failure and bring down the link and the dynamic interface stacks built on the link.

During a unified in-service software upgrade, the LOS does not last more than 2.5 seconds. The remote device detect a momentary LOS but does not perceive this short LOS as a link failure and does not bring the link down,

T3

Local T3 (DS3) devices are reprogrammed during a unified in-service software upgrade, generating a defect. The router completes the reprogramming within 2.5 seconds. Because JUNOS DS3 applications declare an alarm and bring down the link only if the defect persists for more than 2.5 seconds, unified ISSU does not cause the links to be brought down. However, the remote T3 devices must also wait 2.5 seconds before declaring an alarm. If the equipment on the far end of the T3 connection generates an alarm immediately rather than waiting, the link goes down, causing the higher layers to also go down for the remote equipment.

TACACS+ Services Not Available

During the upgrade phase of unified ISSU, TACACS+ services are unavailable. If you have configured AAA authentication for Telnet (with the **aaa new-model command**) this lack of availability affects CLI authentication, authorization, and accounting activities.

CLI login and privilege authentication cannot succeed during a unified ISSU unless you configure at least one of the alternate authentication methods with the **aaa authentication login** command: **enable**, **line**, or **none**.

Similarly, CLI exec and command authorization cannot succeed during a unified ISSU unless you configure one of the alternate authorization methods with the **aaa authorization** command: **if-authenticated** or **none**.

Because there is no alternate method of accounting other than TACACS, CLI exec and command accounting does not work during this phase.

Interruption in Traffic Forwarding for Layer 3 Routing and Signaling Protocols

The routing protocols are affected by two interruptions in traffic forwarding caused by the unified in-service software upgrade during the upgrade phase.

- Switchover from active to standby SRP module—When the active SRP module running the current release switches over to the standby SRP module running the upgrade release, the routing protocols and all other applications restart. A control plane outage of 30–40 seconds prevents the protocols from sending hellos or keepalive messages.

The protocols must gracefully restart to come back online, recover their routing state on the newly active SRP module, and respond to their peers. Therefore, you must enable graceful restart for all protocols before you begin the unified in-service software upgrade. All neighbors of the routing protocols must also be configured to support graceful restart.

A data plane outage of about 1 second for the E120 and E320 routers and about 4 seconds for the ERX1440 router also takes place during the switchover of the fabric from the active primary SRP module to the standby SRP module.

- Upgrade of the forwarding plane for each line module—After the routing protocols reconverge with their peers and rebuild their routing tables, unified ISSU upgrades the forwarding plane on all line modules simultaneously. This upgrade halts forwarding through the chassis. This forwarding outage lasts about 15 seconds for the E120 and E320 routers and about 50 seconds for the ERX1440 router.

If capable, routing protocols temporarily lengthen their timers to survive the outages. During the initialization phase, unified ISSU checks for timers that are set too short and whether the protocol enables timer renegotiation. If these checks fail, unified ISSU generates a warning and enables you to reconfigure the protocols before you issue the **issu start** command.

We recommend that you configure timers to be longer than usual for the routing protocols that cannot renegotiate timers. You can use bidirectional forwarding detection (BFD) to quickly detect forwarding interruptions.

Table 9 on page 87 describes how individual routing protocols behave during a unified in-service software upgrade.

Table 9: Behavior of Routing Protocols During a Unified In-Service Software Upgrade

Protocol	Behavior
BFD	BFD renegotiates its timers as needed. Typically, the timers are lengthened until the SRP module switchover takes place, then shortened for the forwarding plane upgrade, and finally shortened to the original configured values.
BGP	<p>The configured BGP timers are typically long enough to survive the forwarding outages. If not, unified ISSU generates a warning message with a recommended timer interval.</p> <p>BGP sends out keepalive messages immediately before and immediately after both the SRP module switchover and the forwarding plane restart, independent of the interval since it last sent them.</p>
IS-IS	If necessary, temporarily lengthens the hello timers.
LDP	<p>Unified ISSU warns you if the hello timers or the keepalive timers are not long enough to survive the forwarding plane upgrade.</p> <p>LDP sends out hello messages and keepalive messages immediately before and immediately after both the SRP module switchover and the forwarding plane restart, independent of the interval since it last sent them.</p>
OSPF	<p>OSPF timers are not negotiable between peers. Unified ISSU generates a warning if the hello timers or the keepalive timers are not long enough to survive the forwarding plane upgrade.</p> <p>OSPF begins a graceful restart before the SRP module switchover. When you configure graceful restart before the unified in-service software upgrade, you must ensure that the graceful restart times are long enough to allow recovery.</p> <p>OSPF sends out hello messages and keepalive messages immediately before and immediately after forwarding plane restart, independent of the interval since it last sent them.</p>
PIM	<p>If necessary, temporarily lengthens the hold times in hello messages. PIM guarantees that at least one hello message with a lengthened hold time is sent to each neighbor.</p> <p>If necessary, increases the join-prune hold time. PIM guarantees that at least one join-prune message with a lengthened hold time is sent to each neighbor.</p>
RIP	RIP timers do not affect unified ISSU.
RSVP-TE	<p>If necessary, temporarily lengthens the graceful restart timers to survive the SRP module switchover.</p> <p>If necessary, lengthens the hello timers to survive the forwarding plane upgrade.</p>

You might want some or all traffic to be routed around the upgrading router rather than accept a forwarding loss during the forwarding interruption. To do so, you must

configure your routing protocols appropriately. For example, you might raise the link cost in IS-IS and OSPF, causing their neighbors to seek alternate routes that have lower link costs. In PIM, you can set the priority for the router interface to zero to ensure that the upgrading router is not selected as a designated router.

Recommended Routing Protocol Timer Settings

You can use the default values for many of the routing protocol timers with no adverse effect on a unified in-service software upgrade. For other timers, we recommend particular values, as described in Table 10 on page 88.

Table 10: Recommended Routing Protocol Timer Settings

Protocol	Timers
BFD	Use the default timers.
BGP	<p>Use the default timers, including graceful restart default timers. If the expected forwarding outage for the platform is beyond what the BGP session's graceful restart mechanism can survive, the unified ISSU initialization process generates a warning message accordingly. In this event, adjust the timer intervals as advised by the message.</p> <p>For information about the expected forwarding outage, see "Interruption in Traffic Forwarding for Layer 3 Routing and Signaling Protocols" on page 86.</p>
DVMRP	Use the default timers.
IGMP	Use the default timers.
IS-IS	Use the default timers, including graceful restart default timers.
LDP	<p>Use the default timers, including graceful restart default timers, except for the following:</p> <ul style="list-style-type: none"> ■ Set the hello hold time to at least 901 seconds for a helper or a restarter configuration for a link-level adjacency or for LDP targeted sessions.
OSPF	<p>Use the default timers, including graceful restart default timers, except for the dead interval.</p> <p>If the expected forwarding outage for the platform is longer than the configured dead interval, the unified ISSU initialization process generates a warning message accordingly. In this event, adjust the timer interval as advised by the message.</p> <p>For information about the expected forwarding outage, see "Interruption in Traffic Forwarding for Layer 3 Routing and Signaling Protocols" on page 86.</p>

Table 10: Recommended Routing Protocol Timer Settings *(continued)*

Protocol	Timers
PIM	<p>Set the query interval to at least 210 seconds.</p> <p>Unified ISSU generates a warning for any of the following conditions, but you can ignore the warning without causing a higher FC outage:</p> <ul style="list-style-type: none"> ■ The current router is a DR. ■ The current router is configured as an Auto RP mapping agent and is chosen as the RP for any group. ■ The current router is an elected or candidate BSR, or if BSR candidate RPs are configured. ■ The graceful restart timer is less than the default value, 30 seconds.
RIP	<p>Use the default timers; graceful restart is not supported. For scaled configurations, such as for 2000 RIP interfaces, use the following values:</p> <ul style="list-style-type: none"> ■ Flush interval: 600 seconds ■ Holddown time: 260 seconds ■ Invalid interval: 260 seconds ■ Update interval: 60 seconds
RSVP-TE	<p>Use the default timers, including graceful restart default timers, except for the following:</p> <ul style="list-style-type: none"> ■ For graceful restart, the hello timeout interval is the product of hello misses multiplied by the hello refresh interval. Determine which period is longer, the IC upgrade time or the forwarding upgrade time. Configure the hello refresh and hello miss values so that the hello timeout is greater than the longer of those two periods. ■ For node hellos, the product of the refresh misses multiplied by the hello refresh interval must be great than the FC outage time. For an outage time of less than 30 seconds, for example, configure the following values: <ul style="list-style-type: none"> ■ Set the node hello refresh interval to 8000. ■ Set the node hello refresh misses to 4.

Before You Begin a Unified In-Service Software Upgrade

The following hardware and software prerequisites must be met for the successful completion of unified ISSU. You can issue the **show issu** command to determine whether the routers meets these requirements.

Hardware Requirements for Unified ISSU

- The router must support unified ISSU. Therefore the router must be an E120, E320, or ERX1440 router.
- Two SRP modules must be installed in the router.

- All installed combinations of line modules and IOAs must support unified ISSU. Unsupported modules that are online are reloaded during the unified ISSU, with consequent loss of connections and traffic forwarding.

Do not install IOAs in the chassis while the unified ISSU operation is in process.

- The redundant SRP module must have at least 300 MB of free memory. Depending on their configuration, line modules require up to 75 MB of free memory.

For information about modules supported on E120 and E320 routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support unified ISSU.

On the ERX1440 router, certain hardware updates might require a module to be cold restarted. Unified ISSU cannot be successfully accomplished with such modules. In this case, the behavior is the same as for unsupported line modules. The unified ISSU process reboots these modules and holds them down until the supported modules on the router complete the unified ISSU process.

When hardware updates are required for modules that you have installed in an ERX1440 router, contact your Juniper Networks representative to determine whether the update affects unified ISSU.

Software Requirements for Unified ISSU

- The running JUNOS Software release must support unified ISSU.

You can upgrade to a software version that supports unified ISSU from a software version that does not support unified ISSU only by means of a conventional upgrade. During the conventional upgrade, all line modules are reloaded, all subscribers are dropped, and traffic forwarding is interrupted until the upgrade is completed.

- The armed (upgrade) release must be capable of being upgraded to from the currently running release; it must be higher-numbered than the running release.
- All applications that are configured on the router must support unified ISSU and stateful SRP switchover.

If one or more unified ISSU-challenged applications are configured and you proceed with a unified in-service software upgrade, the unified ISSU process forces a conventional upgrade on the router. All line modules are reloaded, all subscribers are dropped, and traffic forwarding is interrupted until the upgrade is completed.

You can avoid this circumstance by removing the configuration for the unified ISSU-challenged applications from the router before you begin the in-service software upgrade.

See “Application Support for Unified ISSU” on page 65 for information about whether an application supports unified ISSU.

- Stateful SRP switchover must be configured on the router. Use the following commands to configure high availability:

```
host1(config)#redundancy
host1(config-redundancy)#mode high-availability
```

See “Managing Module Redundancy” on page 7 for information about high availability.

The following requirements must be met for traffic forwarding to continue. However, failing to meet these requirements does not halt the unified ISSU operation. The unified ISSU process offers the option to override or ignore these forwarding requirements.

- Graceful restart must be enabled for all configured routing protocols. The unified ISSU operation relies on graceful restart to keep the routing protocols alive through the various stages of the upgrade.
- All connected peers must be configured with graceful restart. Because some protocols cannot themselves confirm peer configuration for graceful restart, you must ensure that the peers are properly configured.
- For applications that exchange keepalive messages with peers, you must ensure that the poll times are adequate to maintain the peering session across any forwarding interruption caused by the unified ISSU operation.
- On the ERX1440 router, you must enter the key provided with your license in order to make the unified ISSU CLI commands available. Unified ISSU is licensed on only the ERX1440 router; no license is required or available on the E120 and E320 routers.

The **license issu** command is available only on the ERX1440 CLI.

Upgrading Router Software with Unified ISSU

To upgrade your router software by means of unified ISSU, perform the following steps.

1. Disable autosynchronization.

```
host1(config)#disable-autosync
```

2. Copy the new release to the router.



NOTE: Be sure to specify the correct software release (.rel) filename for the router you are using, as described in the section *Identifying the Software Release File* in the *JUNOS System Basics Configuration Guide*.

```
host1#copy /incoming/releases/ftpserver/quebec2.rel R2.rel
```

3. Save the current configuration.

```
host1#copy running-configuration system2.cnf
```

4. Determine whether the router hardware and the software release meet the criteria required for unified ISSU to operate successfully by using one of the following commands:

```
host1#show issu
host1#show issu brief
host1#show issu detail
```

5. Arm the primary SRP module with the upgrade release.

```
host1#boot system R2.rel
```



NOTE: You must arm any hotfixes that need to be loaded with the new release after you have armed the new release. The hotfixes are supplied when the modules to which they apply are rebooted.

6. Synchronize the NVS file system of the redundant SRP module with that of the primary SRP module.

```
host1#synchronize
```

Because the redundant SRP module is running a different release than the armed release, the module automatically reboots and runs the armed (upgrade) release, R2.rel.

Wait for the redundant SRP module to boot, initialize, and reach the standby state. At this point, the REDUNDANT LED on the module is illuminated and the ONLINE LED is off. The State field in the **show version** display indicates that the redundant module is in the standby state.

7. Synchronize the file system of the primary module with that of the redundant module.

The NVS file systems of the two SRP modules are unsynchronized because the redundant SRP module rebooted.

```
host1#synchronize
```

8. Reenable autosynchronization.

```
host1(config)#no disable-autosync
```

9. (ERX1440 only) Configure the ERX1440 license key.

```
host(config)#license issu xyz123abc
License for ISSU configured.
```

10. Determine whether unified ISSU is in the Idle state and whether all upgrade requirements have been met.

```
host1#show issu
```



NOTE: If the results indicate that some requirements are not met, you must correct this situation before proceeding.

11. Ensure that stateful SRP switchover is configured on the router.

```
host1#show redundancy srp
```

If it is not already configured, do so now.

```
host1(config)#redundancy
host1(config-redundancy)#mode high-availability
```

12. For each configured protocol on the router and its neighbors, ensure that graceful restart is configured. See the relevant protocol configuration chapters in the JUNOS document set for information about configuring graceful restart.
13. Begin the initialization phase of the unified in-service software upgrade.

```
host1#issu initialize
```

The CLI displays the status of the initialization as it proceeds.

14. (Optional) From a different CLI session, display the progress of the initialization.

```
host1#show issu
```

Unified ISSU must be in the Initialized state before you proceed to the next step. The time required for initialization varies with the system load and the complexity of the router configuration.

15. Start the upgrade phase.

```
host1#issu start
```

The router switches to the redundant SRP module running the upgrade release, R2.rel. Significant upgrade milestones are displayed as they occur.

16. When the console indicates that the upgrade is completed, you can verify that the router is back in the idle state and running the upgrade release, R2.rel.

```
host1#show issu
```

You can also verify the status of the SRP modules and line modules, as well as the running and armed releases.

```
host1#show version
```

issu initialize

- Use to start the initialization phase of the unified ISSU process.
- This command displays the percentage completion for the process as it takes place.
- Example

```
host1#issu initialize
Verifying the ISSU criteria... verified
Starting the ISSU initializing phase
Upgrading the standby SRP- This phase can take a long time
10% completed...
```

- There is no **no** version.
- See `issu initialize`.

issu start

- Use to start the upgrade phase of the unified ISSU process after the initialization phase has completed.
- Example

```
host1#issu start
Verifying the ISSU criteria... verified
The system will now enter the upgrading phase. This phase cannot be aborted.
Do you wish to continue?
Yes
Starting the ISSU upgrade phase
... Upgrading the line card – Control plane
... Upgrading completed
Switching from primary SRP to the standby SRP
The system will resume on the SRP in slot 7 in a few minutes.
```

- There is no **no** version.
- See `issu start`.

issu stop

- Use to gracefully stop a unified in-service software upgrade and place the process in an idle state.
- You can issue this command only when unified ISSU is in the initialized state. You cannot issue this command after you have issued the **issu start** command to begin the upgrade phase of unified ISSU.
- Example

```
host1#issu stop
The command will abort the ISSU operation. Do you wish to continue?
Yes
Stopping the ISSU upgrade process
...reloading standby SRP
```

- There is no **no** version.
- See `issu stop`.

license issu

- Use to specify the unified ISSU license key for an ERX1440 router.
- Purchase a unified ISSU license to allow a unified ISSU upgrade on an E Series router.



NOTE: Acquire the ERX1440 unified ISSU license from Juniper Networks Customer Service or your Juniper Networks sales representative.

- The **issu initialization**, **issu start**, **issu stop**, and **show issu** commands are not visible on an ERX1440 router until you specify the license.
- This command is not visible on an E120 or E320 router. A license is not required or available for the E120 and E320 routers.
- Example 1—the correct key is entered

```
host1#license issu xyz123abc
License for ISSU configured.
```

- Example 2—an incorrect key is entered

```
host1#license issu abc123def456
%Unable to configure ISSU license (Invalid License)
```

- Use the **no** version to disable the license on the ERX1440 router.
- See `license issu`

Halting the Unified ISSU Process and Restoring the Original State of the Router

The options that are available to halt the unified in-service software upgrade depend on the phase that the upgrade is in when you attempt to halt it. The phase also affects the state of the router after the upgrade is halted.

Halting Unified ISSU During Initialization Phase

During the initialization phase, you can halt the unified ISSU process by issuing the **issu stop** command. This action reloads the redundant SRP module with the armed upgrade release. As a result, unified ISSU is placed in the idle state and the following releases are present on the router:

- Primary SRP module—Running (original) release
- Redundant SRP module—Upgrade release
- Line modules—Running (original) release

After you stop unified ISSU, you can return the router to the state it was in when you began the unified in-service software upgrade. To roll the router back to its beginning state with the redundant SRP module running the original release, you must perform the following steps to arm the redundant SRP module with the running release:

1. Turn off auto synchronization.

```
host1(config)#disable-autosync
```

2. Specify that the router use the running release when it reboots.

```
host1(config)#boot system erx_x-y-z.rel
```

3. Synchronize the NVS file system of the redundant module with that of the primary module.

```
host1#synchronize
```

The redundant SRP module automatically reboots because the armed release (the original release) now differs from the software release it is currently running (the upgrade release).

4. Verify that stateful SRP switchover is enabled.

```
host1#show redundancy
```

Halting Unified ISSU During Upgrade Phase

During the upgrade phase—before the line module and control plane software is upgraded—the unified ISSU process provides an opportunity to cancel the upgrade. If you choose to cancel, the router remains in the unified ISSU initialized state. The CLI command set becomes fully accessible.

If you do not cancel at this point, then the process continues and any line modules that do not support unified ISSU are reloaded. Application sessions are brought down and traffic forwarding is interrupted for the unsupported modules.

If you do cancel in response to the CLI prompt, unified ISSU returns to the initialized state, and the following releases are present on the router:

- Primary SRP module—Running (original) release
- Redundant SRP module—Upgrade release; the module is in the unified ISSU initialized state
- Line modules—Running (original) release

To roll back from the unified ISSU initialized state, you must issue the **issu stop** command. The command reloads the redundant SRP module with the armed release and places unified ISSU in the idle state. As a result, the following releases are present on the router:

- Primary SRP module—Running (original) release
- Redundant SRP—Upgrade release

- Line modules—Running (original) release

After you stop unified ISSU, you can return the router to the state it was in when you began the unified in-service software upgrade. To roll the router back to its beginning state with the redundant SRP module running the original release, you must perform the following steps to arm the redundant SRP module with the running release:

1. Turn off auto synchronization.

```
host1(config)#disable-autosync
```

2. Specify that the router use the running release when it reboots.

```
host1(config)#boot system erx_x-y-z.rel
```

3. Synchronize the NVS file system of the redundant module with that of the primary module.

```
host1#synchronize
```

The redundant SRP module automatically reboots because the software release that it is configured to run now differs from the software release it is running.

Monitoring a Unified In-Service Software Upgrade

You can use the **show issu** command to monitor the status of the router with regard to a unified in-service software upgrade.

show issu

- Use to display information about the current status of the router relative to a unified in-service software upgrade and of the upgrade itself.
- Field descriptions
 - ISSU state—State of the upgrade process, idle, initializing, initialized, or upgrading
 - ISSU description—State of the upgrade, including percent complete
 - criteria met—Whether prerequisites for the upgrade have been met and, generally, what errors occurred
 - running release—Filename of JUNOS Software release that is currently running on the SRP modules
 - armed release—Filename of JUNOS Software release that is armed to become the next running release when the router reboots
- Example 1—Displays the current unified ISSU state and identifies the active and armed releases

```
host1#show issu brief
```

```
ISSU state:      initializing
ISSU description: ISSU initialize is in-progress, 5% complete
criteria met:    No, upgrade error(s) found
running release: release1.rel
armed release:   release2.rel
```

- Example 2—To the information displayed by **show issu brief**, adds a summary table of unified ISSU verification criteria

```
host1# show issu
ISSU state:      idle
ISSU description: ISSU is currently idle
criteria met:    No, upgrade error(s) found
running release: release1.rel
armed release:   release2.rel
#               ISSU Activation Criteria Summary               Met
```

--	-----	-----
1	In-Service Software Upgrade ready?	Yes
2	High-Availability ready?	No
3	Line modules ready?	Conditional
4	Configuration conversion support ready?	Yes
5	CLI sessions ready?	Yes
6	Routing applications ready?	Yes
7	Protocol timers ready?	Yes

- Example 3—To the information displayed by **show issu**, adds a detailed table of unified ISSU verification criteria that lists mandatory and conditional criteria that have not been met, the impact of this status, and the remedy as reported by router applications and system components that participate in the in-service software upgrade

```
host1# show issu detail
ISSU state:      idle
ISSU description: ISSU is currently idle
criteria met:    No, upgrade error(s) found
running release: release1.rel
armed release:   release2.rel
#               ISSU Activation Criteria Summary               Met
```

--	-----	-----
1	In-Service Software Upgrade ready?	Yes
2	High-Availability ready?	No
3	Line modules ready?	Conditional
4	Configuration conversion support ready?	Yes
5	CLI sessions ready?	Yes
6	Routing applications ready?	Yes
7	Protocol timers ready?	Yes
#	ISSU Criterion Detail	Met
--	-----	-----


```

1   In-Service Software Upgrade ready?                Yes
2   High-Availability ready?                          No
->  Problem: The standby SRP must not be running the same release No

      Reporting Slot: 6
      Impact: ISSU cannot be performed
      Remedy: boot a release compatible with ISSU on the standby SRP

3   Line modules ready?                               Conditional
->  Problem: Card does not support required memory configuration Conditional
      : Slot 1, OC3/OC12/DS3-ATM, requires at least 256 MB
      Reporting Slot: 6
      Impact: If you continue, the card will immediately be reset
      and then cold started when ISSU Upgrade completes
      Remedy: data unavailable
->  Problem: Card does not support required memory configuration Conditional
      : Slot 8, CT3-12, requires at least 256 MB
      Reporting Slot: 6
      Impact: If you continue, the card will immediately be reset
      and then cold started when ISSU Upgrade completes
      Remedy: data unavailable
->  Problem: Card does not support required memory configuration Conditional
      : Slot 9, CT3-12, requires at least 256 MB
      Reporting Slot: 6
      Impact: If you continue, the card will immediately be reset
      and then cold started when ISSU Upgrade completes
      Remedy: data unavailable
->  Problem: Card does not support required memory configuration Conditional
      : Slot 10, CT3-12, requires at least 256 MB
      Reporting Slot: 6
      Impact: If you continue, the card will immediately be reset
      and then cold started when ISSU Upgrade completes
      Remedy: data unavailable
->  Problem: Card not disabled or not online: Slot 1, OC3/OC12/D Conditional
      S3-ATM, 0/0
      Reporting Slot: 6
      Impact: If you continue, the card will immediately be reset
      and then cold started when ISSU Upgrade completes
      Remedy: If not standby, Wait for card to come online before
      proceeding
->  Problem: Card not disabled or not online: Slot 8, CT3-12, 0/ Conditional
      0
      Reporting Slot: 6
      Impact: If you continue, the card will immediately be reset
      and then cold started when ISSU Upgrade completes
      Remedy: If not standby, Wait for card to come online before
      proceeding
4   Configuration conversion support ready?           Yes
5   CLI sessions ready?                              Yes
6   Routing applications ready?                      Yes
7   Protocol timers ready?                           Yes

```

■ See show issu.

Chapter 5

Configuring VRRP

This chapter describes how to configure the Virtual Router Redundancy Protocol (VRRP) on your E Series router.

- VRRP Overview on page 101
- Platform Considerations on page 102
- References on page 102
- How VRRP Works on page 103
- How VRRP Is Implemented in E Series Routers on page 106
- Configuring VRRP on page 108
- Changing Object Priority on page 112
- Monitoring VRRP on page 113

VRRP Overview

VRRP can prevent loss of network connectivity to end hosts if the static default IP gateway fails. By implementing VRRP, you can designate a number of routers as *backup* routers in the event that the default *master* router fails. VRRP fully supports Virtual Local Area Networks (VLANs) and stacked VLANs (S-VLANs).



NOTE: The term *virtual router* as defined in *Configuring Virtual Routers* in the *JUNOS System Basics Configuration Guide*, is different from what is implied by VRRP. In this chapter, the term *virtual router* always refers to a VRRP router; that is, a router that has enabled VRRP.

In case of a failure, VRRP dynamically shifts the packet-forwarding responsibility to a backup router. VRRP creates a redundancy scheme that enables hosts to keep a single IP address for the default gateway but maps the IP address to a well-known virtual MAC address. VRRP provides this redundancy without user intervention or additional configuration at the end hosts.

VRRP Terms

Table 11 on page 102 provides definitions for the basic VRRP terms used in this chapter.

Table 11: VRRP Definitions

Term	Definition
VRRP router	<p>A router that is running VRRP. It might participate in one or more virtual router IDs (VRIDs). An IP redundancy instance can:</p> <ul style="list-style-type: none"> ■ Act as a master with associated addresses it owns at an IP interface ■ Act simultaneously as a backup for other routers with additional VRID mappings and priorities for those routers
Master router	The VRRP router that takes the responsibility of forwarding packets sent to the IP addresses associated with the virtual router, and that answers ARP requests for these IP addresses. If the IP address owner is available, it always becomes the master.
Backup router	The VRRP router available to take forwarding responsibility if the current master router fails.
IP address owner	The IP interface–VRID pair instance that has the associated IP addresses as real interface addresses. This router, when up, responds to packets addressed to one of these IP addresses for Internet Control Message Protocol (ICMP) pings or Transmission Control Protocol (TCP) connections. The IP address owner is the <i>primary router</i> .
Primary IP address	An IP address configured as primary from the set of real interface addresses. VRRP advertisements are always sent (by the master router) using the primary IP address as the source of the IP packet.

Platform Considerations

For information about modules that support VRRP on ERX14xx models, ERX7xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support VRRP.

For information about modules that support VRRP on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support VRRP.

References

For more information about VRRP, see:

- RFC 2338—Virtual Router Redundancy Protocol (April 1998)

- RFC 2787—Definitions of Managed Objects for the Virtual Router Redundancy Protocol (March 2000)



NOTE: We recommend that you have some background understanding of the Address Resolution Protocol (ARP) before you configure VRRP. See *Address Resolution Protocol* in the *JUNOS IP, IPv6, and IGP Configuration Guide*.

How VRRP Works

The advantage of using VRRP is that you gain a higher availability for the default path without requiring configuration of dynamic routing or router discovery protocols on every end host.

VRRP routers viewed as a *redundancy group* share the responsibility for forwarding packets as if they *owned* the IP address corresponding to the default gateway configured on the hosts. At any time, one of the VRRP routers acts as the master, and other VRRP routers act as backup routers. If the master router fails, a backup router becomes the new master. In this way, router redundancy is always provided, allowing traffic on the LAN to be routed without relying on a single router.

A master always exists for the shared IP address. If the master goes down, the remaining VRRP routers elect a new master VRRP router. The new master forwards packets on behalf of the owner by taking over the virtual MAC address used by the owner.

When implemented in your network, VRRP interprets any active link to a subnet to indicate the router has access to the entire subnet. VRRP leverages the broadcast capabilities of Ethernet. Provided that one of the routers in a VRRP configuration is running, ARP requests for the IP addresses assigned to the default gateway always receive replies. Additionally, end hosts can send packets outside their subnet without interruption.

Configuration Examples

This section describes and illustrates three VRRP configuration examples.

- “Basic VRRP Configuration” on page 103
- “Commonly Used VRRP Configuration” on page 104
- “VRRP Configuration Without the Real Address Owner” on page 105

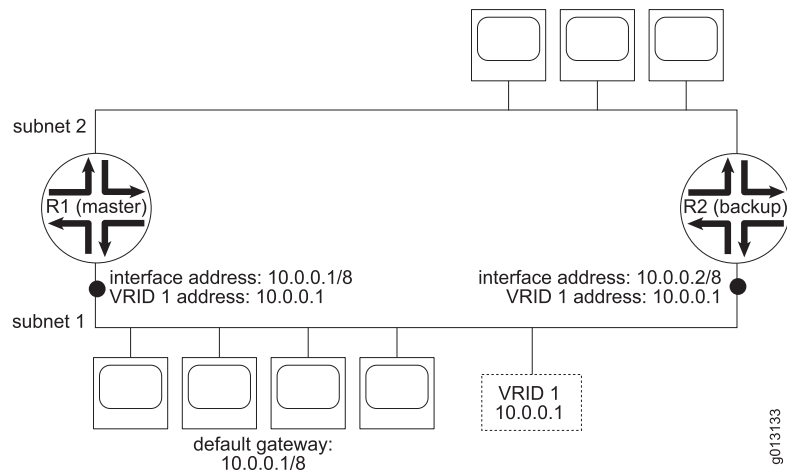
Basic VRRP Configuration

As Figure 4 on page 104 shows, the basic VRRP configuration uses a single VRID (VRID 1). Because R1 is the address owner, it serves as the master router. Router R2 is the backup router. The four end hosts on subnet 1 are configured to use 10.0.0.1/8 as the default router. IP address 10.0.0.1 is associated with VRID 1.

In this example, if R1 becomes unavailable, R2 takes over VRID 1 and its associated IP addresses. Packets sent to IP destinations outside the 10.x.x.x subnet using 10.0.0.1 as the router are then forwarded by R2. Even though R2 assumes R1's forwarding responsibilities, it may or may not process any packet with destination address (DA) 10.0.0.1, depending on the accept-data configuration. When R1 becomes active again, it takes over as the master router and R2 reverts to the backup router.

The VRRP MAC address is always 00-00-5e-00-01-vrid. The valid VRID range is 0x01–0xFF.

Figure 4: Basic VRRP Configuration



Commonly Used VRRP Configuration

Figure 5 on page 105 shows two physical routers backing up each other through VRRP. Routers R1 and R2 are both configured with VRID 1 and VRID 2. In this configuration, under normal circumstances the routing load is distributed between the two routers.



Figure 5: Commonly Used VRRP Configuration**VRRP Configuration Without the Real Address Owner**

Figure 6 on page 106 is noticeably similar to Figure 5 on page 105 except that the addresses configured by the VRIDs have no real owner. Consequently, both routers R1 and R2 are configured as backup routers for VRID 1, VRID 2, VRID 3, and VRID 4.

Figure 6: VRRP Configuration Without the Real Address Owner

Assuming that preemption is enabled, the router that is configured with the highest priority for each VRID becomes the master router. If priorities are the same, the router that has the highest primary address becomes the master router.

This configuration shows how the address owner does not necessarily need to exist under VRRP, and all PCs can reach destinations outside of their network through the current master VRRP router. Depending on the accept-data configuration, the PCs may even be able to ping their default gateway.

The election protocol specified in VRRP uses IP multicast packets to provide the router with redundancy. Therefore, VRRP can operate over a variety of multiaccess LAN technologies that support IP multicast. It is important to remember that there is always one master router for an IP address shared by the redundancy group.

How VRRP Is Implemented in E Series Routers

VRRP is implemented in E Series routers to meet two goals. The first goal is to avoid the single point of failure inherent to hosts that have a single default gateway configured. The second goal is to keep the complexity of redundancy away from the hosts themselves. These goals comply with RFC 2338 and RFC 2787.

The association between VRIDs and IP addresses is coordinated among all participating VRRP routers. The following scenario can help you understand how VRRP is implemented in the router.

1. An E Series router assigns common VRIDs to the group of routers that are going to share IP addresses.
2. The E Series router sends VRRP advertisements to well-known multicast addresses. The router that owns the addresses automatically becomes the master and sends periodic VRRP advertisement messages. A VRRP advertisement consists of the IP addresses that the master router controls and the VRID.

3. If the master router stops advertising for a predetermined period of time, the remaining routers using the same VRID enter an election process to determine which router takes over the master router responsibilities.
4. Depending on the configuration, the master router that does not own the IP addresses might do one of the following:
 - Drop all packets that have destination addresses to these IP addresses (default)
 - Accept packets that have destination addresses to these IP addresses as if the addresses belonged to the master router (using the **ip vrrp accept-data** command).
5. If the elected master router fails, backup routers start the election process again.
6. When the original master router becomes operational again, it restarts broadcasting advertisements as long as preemption is enabled or the master router is the address owner. Packet forwarding responsibility then shifts back to the original master router.

Router Election Rules

If the master router becomes unavailable, the following rules govern election of the master router:

- The backup router assigned the highest priority for each VRID becomes the master router.
- If two backup routers were assigned the same priority, the router that has the highest primary address becomes the master router. For example, if several routers were all assigned the default priority of 100, the IP addresses must be compared.
- Router election on a VRRP router can also be determined by whether the preemption option is enabled.

When a backup router detects a master router with a lower priority than the backup router has, the backup router might leave the current master router alone or take over the current master router and become the master router itself.

When preemption is enabled, a backup router always preempts or takes over the responsibility of the master router. When preemption is disabled, the lower-priority backup is left in the master state.



NOTE: Using VRRP can override the source address of the ICMP redirect. When a backup VRID functions as a master router on a given IP interface, its ICMP redirects must *fake* the source IP address of the IP address owner. The redirect must fake the IP address because hosts accept only an ICMP redirect that is sent by the current gateway of the host.

Configuring VRRP

Configuring VRRP requires that you first configure an IP interface over which you can configure VRRP and any VRID instances in which you want the VRRP routers to participate. The following sections contain information for configuring the IP interface for VRRP, any VRID instances for the VRRP routers, and steps for creating a basic VRRP configuration.

Configuring the IP Interface

Before you configure VRRP, you must configure an IP interface and assign a primary IP address and subnet mask. When the IP address belongs to the owner of the VRID, you must associate the IP address with the VRID that you create.

To configure the IP interface for VRRP:

1. Configure an IP interface.

```
host1(config)#interface fastEthernet 4/0
```

2. Assign an IP address and a subnet mask.

```
host1(config-if)#ip address 194.50.1.42 255.255.255.0
```



NOTE: We recommend that you complete all IP address configurations before you configure VRRP. If for any reason the IP address information changes after you configure VRRP, you must revise the associated IP addresses configured on the related VRRP entries. If you specify **auto** addresses in the **ip vrrp virtual-address** command along with using priority 255, you must disable and reenab the VRRP entry to update the association list.

Creating VRIDs

A master or backup router running the VRRP protocol can participate in one or more VRID instances. You can create a VRID instance in several ways:

- We recommend that you create and configure a VRID instance first, and enable it last. For example:

```
host1(config-if)#ip vrrp 198  
host1(config-if)#ip vrrp 198 priority 255
```

- You can create and enable a VRID instance by using the **ip vrrp vrid enable** command. For example:

```
host1(config-if)#ip vrrp 25 enable
```

- You continue to configure the VRID by identifying the VRID each time you use a VRRP command. For example:

```
host1(config-if)#ip vrrp 175 authentication-type none
```

- Alternatively, you can create a new VRID when you use any VRRP command, provided that you are using the VRID instance for the first time. For example, if you execute the **ip vrrp vrid preempt** command and it is the first time that you use the VRID, the command creates a new VRID.

```
host1(config-if)#ip vrrp 16 preempt
```

- Use the **ip vrrp vrid enable** command last. The new VRID is not enabled until you execute this command.

```
host1(config-if)#ip vrrp 198 enable
host1(config-if)#ip vrrp 16 enable
host1(config-if)#ip vrrp 175 enable
```

Configuration Steps

Before you configure VRRP, we recommend that you review the configuration examples in the earlier section “How VRRP Works” on page 103.

To configure VRRP parameters:

1. (Optional) Create a VRID instance.

```
host1(config-if)#ip vrrp 25
```

2. (Optional) Set a VRRP advertisement interval for the same VRID.

```
host1(config-if)#ip vrrp 25 advertise-interval 50
```

3. Set the VRRP router priority for owner or backup routers.

This step is mandatory to configure priority for the owner VRID (255). This step is optional to configure priority for a backup VRID (1–254). The default value is 100.

```
host1(config-if)#ip vrrp 25 priority 255
host1(config-if)#ip vrrp 22 priority 254
```

4. (Optional) Specify that the backup router can process packets with an IP destination address of the virtual address.

```
host1(config-if)#ip vrrp 22 accept-data
```

5. (Optional) Set the preempt option. This example creates a new VRID.

```
host1(config-if)#ip vrrp 10 preempt
```

6. Associate an IP address with a VRID.

```
host1(config-if)#ip vrrp 25 virtual-address 194.2.1.63
```

7. (Optional) Set the VRRP authentication type to either **text** or **none**.

```
host1(config-if)#ip vrrp 25 authentication-type none
```

8. (Optional) Configure the VRRP authentication key.

```
host1(config-if)#ip vrrp 25 authentication-key dublin
```

9. Enable the VRID instance.

```
host1(config-if)#ip vrrp 25 enable
```

ip vrrp

- Use to create a VRID instance.
- The VRID range is 1–255.
- Example

```
host1(config-if)#ip vrrp 25
```

- Use the **no** version to remove a VRID instance.
- See `ip vrrp`.

ip vrrp accept-data

- Use to enable the backup router to process packets with an IP destination address equivalent to the virtual addresses while the backup router is in the master state.
- Use the default state (disabled) for full compliance with RFC 2338.
- The configuration ignores this attribute if the VRRP entry uses a priority of 255 (owner).



NOTE: When using this attribute and also restricting incoming packets to ICMP only, you must use policy filters to accept only ICMP packets with the virtual address as the destination address.

- Example

```
host1(config-if)#ip vrrp 22 accept-data
```

- Use the **no** version to disable processing of data packets by the backup router while the router is in the master state. When disabled, the master router drops any packets with an IP destination address equivalent to the virtual addresses.
- See `ip vrrp accept-data`.

ip vrrp advertise-interval

- Use to configure the amount of time the VRRP router waits between advertisements.
- Specify the interval time in seconds or milliseconds.
- Use seconds to be in compliance with RFC 2338.
- If your VRID environment consists of only E Series routers, you can optionally use milliseconds.
- Example

```
host1(config-if)#ip vrrp 25 advertise-interval 50
```

- Use the **no** version to restore the default value, 1 second.
- See ip vrrp advertise-interval.

ip vrrp authentication-key

- Use to specify the authentication key.
- Use the **key** keyword only when the authentication type is **text**. See the **ip vrrp authentication-type** command.
- Example

```
host1(config-if)#ip vrrp 25 authentication-key dublin
```

- Use the **no** version to set the authentication key to its default, an empty string.
- See ip vrrp authentication-key.

ip vrrp authentication-type

- Use to specify the authentication type; **text** or **none**.
- Example

```
host1(config-if)#ip vrrp 175 authentication-type none
```

- Use the **no** version to set the authentication type to its default, none.
- See ip vrrp authentication-type.

ip vrrp enable

- Use to enable an existing VRID instance.
- Specify a VRID in the range 1–255.
- The default is that VRRP is disabled.
- Example

```
host1(config-if)#ip vrrp 175 enable
```

- Use the **no** version to disable an existing VRID instance.
- See ip vrrp enable.

ip vrrp preempt

- Use to enable preemption. When preemption is enabled, a backup router always takes over the responsibility of the master router. When preemption is disabled, the lower-priority backup router is left in the backup state.

- Example

```
host1(config-if)#ip vrrp 10 preempt
```

- The default is that VRRP preemption is enabled.
- Use the **no** version to disable preemption.
- See ip vrrp preempt.

ip vrrp priority

- Use to configure the priority of VRRP routers.
- Use a value of 255 to imply *master router* priority.
- Use a value in the range 1–254 to imply *backup router* priority.
- Example

```
host1(config-if)#ip vrrp 25 priority 255
```

- Use the **no** version to set the priority to the default value, 100.
- See ip vrrp priority.

ip vrrp virtual-address

- Use to associate an IP address with a VRID.
- Use the **auto** keyword to automatically learn or configure associated addresses, depending on the priority attribute.
- There is no default.
- Example

```
host1(config-if)#ip vrrp 25 virtual-address 194.2.1.63
```

- Use the **no** version to remove an IP address association with a VRID. If you use **auto** addressing, the **no** version clears the **auto** flag.
- See ip vrrp virtual-address.

Changing Object Priority

You can use the **ip vrrp track** command (in conjunction with the **track** command) to track an object by its virtual router ID (VRID). When the state of the object changes from an up state to a down state, the priority of the vrid is decremented. When the object changes back to an up state the priority is restored. You can specify a priority value in the range 1–255 to be used for modifying the priority; the default value is 10.



NOTE: For information about the **track** command, see *Managing the System* in the *JUNOS System Basics Configuration Guide*.

ip vrrp track

- Use to dynamically change the priority of a virtual router ID (VRID) in response to a change in the state of a specified object. You can specify the value by which the priority changes in the range 1–255 or use the default value (10). Multiple VRIDs can track the same object and a single VRID can track multiple objects. The object priority is restored when the state of the object returns to an up state.
- Example 1


```
host1(config-if)#ip vrrp 25 track abc
```
- Example 2


```
host1(config-if)#ip vrrp track xyz decrement 15
```
- Use the **no** version to disable any tracking for the object.
- See ip vrrp track.

Monitoring VRRP

You can use several VRRP show commands to *display* the details of your VRRP configuration.

baseline ip vrrp

- Use to establish the baseline on all VRRP statistics as the current value.
- Example


```
host1#baseline ip vrrp
```
- See baseline ip vrrp.

show ip vrrp

- Use to display a detailed summary of all configured VRIDs.
- Use the **interface** keyword to specify a specific Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface.
- Use the **summary** keyword to display a summary count on all configured VRIDs
- Field descriptions
 - Interface—Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface specifier and VRID
 - primary address—IP address used while in master state; not necessarily an associated address

- operational state—State of the VRRP router: master, backup, or init; when the operational state is backup, the current master router IP address is provided
- admin state—Administrative status: enabled or disabled
- up time—Number of seconds that the VRID has been enabled in non-init state
- interval—VRRP advertisement interval in seconds or milliseconds
- Learning timer mode—Mode of the VRRP router: enabled or disabled; when the mode is enabled, the router learns the VRRP advertisement interval that is useful in case of failure of the master router. If the mode is disabled, the router does not learn the VRRP advertisement interval.
- last error status—Help text used to debug any error detected
- priority—Priority value of VRRP router
- admin priority—Priority of the VRRP administrative router
- auth type—The VRRP authentication type: none or text
- preemption—VRRP router preemption status: enabled or disabled
- accept data—VRRP router accept data status: enabled or disabled
- assoc address(es)—IP addresses associated with the VRID
- track object—Name and state of the tracked object and the value by which the object priority changes following an object state change
- ip interfaces with vrrp—Number of IP interfaces using VRRP
- entries—Total number of entries
- entries enabled—Number of enabled entries
- entries with owner priority—Number of entries with an owner priority
- entries in init state—Number of entries in an initialization state
- entries in backup state—Number of entries in a backup state
- entries in master state—Number of entries in a master state
- entries performing tracking—Number of entries performing tracking functions
- Example 1

```

host1#show ip vrrp
Interface: FastEthernet3/0 vrrpVrid: 1
  primary address: 12.60.1.1
  operational state: init
  admin state: disabled
  up time: N/A
  interval: 1 second
  Learning timer mode: disabled

```



```

last error status: no error
priority: 100 ( admin priority: 100 )
auth type: none
preemption: enabled
accept data: disabled
assoc address(es): none
track object: xyz state: Up decrement: 10

```

■ Example 2

```

host1#show ip vrrp summary
ip interfaces with vrrp: 1
entries: 10
entries enabled: 10
entries with owner priority: 1
entries in init state: 0
entries in backup state: 9
entries in master state: 1
entries performing tracking: 2

```

■ See show ip vrrp.

show ip vrrp brief

- Use to display a brief summary of all configured VRIDs.
- Use the **interface** keyword to specify a specific Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface.
- Field descriptions
 - Interface—Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface specifier
 - VRID—VRRP router instance configured on this interface
 - Primary Address—IP address used while in master state; not necessarily an associated address
 - State—Operational state of VRRP router: master, backup, or init
 - Adv—Advertisement interval, in seconds
 - Pri—Priority assigned to this router
 - Admin—Administrative state of the VRID: enabled or disabled
- Example

```

host1#show ip vrrp brief
Interface          VRID  Primary Address  State  Adv  Pri  Admin
-----
fastEthernet12/8.1.1  255  123.123.123.123  init   1  100  disabled
gigabitEthernet12/8.1.1  1    1.1.1.1         master  1  254  enabled

```

■ See show ip vrrp.

show ip vrrp neighbor

- Use to display neighbors currently known to the VRRP routers.
- A neighbor—a router that shares a given VRID with the VRRP router—is known to the VRRP router only when the neighbor becomes a master for an IP address and sends VRRP advertisements to that effect. If a router sharing the VRID has not yet become a master, then the local router remains unaware of this neighbor and this command does not display that neighbor.
- Use the **interface** keyword to specify a specific Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface.
- Field descriptions
 - Interface—Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface specifier and VRID of neighbors known to the VRRP router
 - time discovered—Date and time that the neighbor was detected
 - primary address—Primary IP address of neighbor
 - adv interval (sec)—VRRP advertisement interval in seconds
 - priority—Priority status of VRRP router (255 = owner)
 - auth type—VRRP authentication type: none or text
 - assoc address(es)—IP addresses associated with the VRID that are advertised by the neighbor
- Example

```

host1#show ip vrrp neighbor
Interface: fastEthernet5/0.0 vrrpVrid: 1
  time discovered: 08/09/2001 07:44
  primary address: 10.0.0.1
  adv interval (sec): 1
  priority: 255 (owner)
  auth type: none
  assoc address(es): 10.0.0.1, 100.0.0.1, 101.0.0.1

Interface: fastEthernet5/0.1 vrrpVrid: 11
  time discovered: 08/09/2001 07:44
  primary address: 11.0.0.1
  adv interval (sec): 1
  priority: 255 (owner)
  auth type: none
  assoc address(es): 11.0.0.1, 110.0.0.1, 111.0.0.1

```

- See show ip vrrp neighbor.

show ip vrrp statistics

- Use to display statistics of configured VRRP routers and each individual VRID.
- Use the **delta** keyword with the **show ip vrrp statistics** command to view the baseline statistics.

- Use the **interface** keyword with the **show ip vrrp statistics** command to specify a specific Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface.
- Field descriptions
 - checksumErrors—Total number of VRRP packets received with an invalid VRRP checksum value
 - versionErrors—Total number of VRRP packets received with an unknown or unsupported version number
 - vridErrors—Total number of VRRP packets received with an invalid VRID for this virtual router
 - iccErrors—Count of line module notifications that did not make it to the controller
 - txErrors—Count of advertisements that did not get sent due to resource limitations
 - rxErrors—Count of advertisements received that could not be parsed by VRRP applications
 - Interface—Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface specifier and VRID
 - becomeMaster—Total number of times that this VRID state has transitioned to master
 - advertiseRcvd—Total number of VRRP advertisements received
 - advertiseIntervalErrors—Total number of VRRP advertisement packets received for which the advertisement interval is different from the one configured for the VRID
 - authFailures—Total number of VRRP packets received that do not pass the authentication check
 - ipTtlErrors—Total number of VRRP packets received with IP TTL (time-to-live) not equal to 255
 - priorityZeroPktsRcvd—Total number of VRRP packets received with a priority of 0
 - priorityZeroPktsSent—Total number of VRRP packets sent with a priority of 0
 - invalidTypePktsRcvd—Total number of VRRP packets received with an invalid value in the Type field
 - addressListErrors—Total number of VRRP packets received for which the address list does not match the locally configured list for the VRID
 - invalidAuthType—Total number of VRRP packets received with an unknown authentication type

- **authTypeMismatch**—Total number of VRRP packets received with an authentication type not equal to the locally configured authentication method
- **packetLengthErrors**—Total number of VRRP packets received with a packet length less than the length of the VRRP header
- **Example 1**—statistics per interface

```
host1#show ip vrrp statistics interface fastEthernet 4/0
```

```
Globals:
```

```
checksumErrors: 0
versionErrors: 0
vrIdErrors: 1
iccErrors: 0
txErrors: 0
rxErrors: 0
```

```
Interface: fastEthernet4/0 vrrpVrid: 1
```

```
becomeMaster: 10
advertiseRcvd: 0
advertiseIntervalErrors: 0
authFailures: 0
ipTtlErrors: 0
priorityZeroPktsRcvd: 0
priorityZeroPktsSent: 9
invalidTypePktsRcvd: 0
addressListErrors: 0
invalidAuthType: 0
authTypeMismatch: 0
packetLengthErrors: 0
```

```
Interface: fastEthernet4/0 vrrpVrid: 50
```

```
becomeMaster: 0
advertiseRcvd: 1000
advertiseIntervalErrors: 0
authFailures: 0
ipTtlErrors: 0
priorityZeroPktsRcvd: 0
priorityZeroPktsSent: 0
invalidTypePktsRcvd: 0
addressListErrors: 0
invalidAuthType: 0
authTypeMismatch: 0
packetLengthErrors: 0
```

- **Example 2**—statistics per interface and VRID

```
host1#show ip vrrp statistics interface fastEthernet 4/0 1
```

```
Interface: fastEthernet4/0 vrrpVrid: 1
```

```
becomeMaster: 0
advertiseRcvd: 0
advertiseIntervalErrors: 0
authFailures: 0
ipTtlErrors: 0
priorityZeroPktsRcvd: 0
priorityZeroPktsSent: 0
invalidTypePktsRcvd: 0
addressListErrors: 0
invalidAuthType: 0
```

```
authTypeMismatch: 0
packetLengthErrors: 0
```

- See show ip vrrp statistics.

show ip vrrp statistics global

- Use to display the statistics of configured VRRP routers and each individual VRID.
- Use the **delta** keyword with the **show ip vrrp statistics global** command to view the baseline statistics.
- Field descriptions
 - checksumErrors—Total number of VRRP packets received with an invalid VRRP checksum value
 - versionErrors—Total number of VRRP packets received with an unknown or unsupported version number
 - vrIdErrors—Total number of VRRP packets received with an invalid VRID for this virtual router
 - iccErrors—Count of line module notifications that did not make it to the controller
 - txErrors—Count of advertisements that did not get sent due to resource limitations
 - rxErrors—Count of advertisements received that could not be parsed by VRRP applications
- Example

```
host1#show ip vrrp statistics global
Globals:
checksumErrors: 0
versionErrors: 0
vrIdErrors: 0
iccErrors: 0
txErrors: 0
rxErrors: 0
```

- See show ip vrrp statistics.

show ip vrrp summary

- Use to display a summary count on all configured VRIDs.
- Field descriptions
 - ip interfaces with vrrp—Total number of VRIDs configured on IP interfaces
 - entries—Total number of entries in all states
 - entries enabled—Number of entries that were enabled
 - entries with owner priority—Number of entries with owner priority

- entries in init state—Number of entries in the init state
- entries in backup state—Number of entries in the backup state
- entries in master state—Number of entries in the master state
- Example

```
host1#show ip vrrp summary
ip interfaces with vrrp: 1
entries: 10
entries enabled: 10
entries with owner priority: 1
entries in init state: 0
entries in backup state: 9
entries in master state: 1
```

- See show ip vrrp.

show ip vrrp tracked-objects

- Use to display details of objects tracked by various VRIDs
- Field descriptions
 - Interface—Name of the interface
 - Vrid—VRRP router instance configured on the interface
 - Priority—Priority of the VRRP router
 - Object—Name of the object being tracked
 - Type—Type of object being tracked
 - State—State of the object
 - Decrement—Value by which the priority is decremented or restored following an object state change
- Example

```
host1#show ip vrrp tracked-objects
```

Interface	Vrid	Priority	Object	Type	State	Decrement
FastEthernet3/0	1	100	ERX_Bangalore	IP-route	Up	12
FastEthernet3/0	1	100	ERX_Bangalore	IP-route	Up	15
FastEthernet3/0	1	100	ERX_Bangalore	IP-route	Up	10
FastEthernet3/0	2	100	ERX_Bangalore	IP-route	Up	10
FastEthernet3/0	3	100	ERX_Bangalore	IP-route	Up	12
FastEthernet3/0	3	100	ERX_Bangalore	IP-route	Up	15

- See show ip vrrp tracked-objects.

Chapter 6

Managing Interchassis Redundancy

This chapter describes how to configure interchassis redundancy (ICR) on your E Series router.

- ICR Overview on page 121
- ICR Platform Considerations on page 123
- ICR Terms on page 124
- ICR References on page 124
- ICR Scaling Considerations on page 124
- Guidelines for Deploying an ICR Partition in Your Network on page 126
- Interaction with RADIUS for ICR on page 127
- Configuring an ICR Partition on page 129
- Configuring the Interface on Which the ICR Partition Resides on page 130
- Configuring VRRP Instances to Match ICR Requirements on page 131
- Naming ICR Partitions on page 131
- Grouping ICR Subscribers Based on S-VLAN IDs on page 132
- Grouping ICR Subscribers Based on VLAN IDs on page 133
- Example: Configuring ICR Partitions That Group Subscribers by S-VLAN ID on page 134
- Using RADIUS to Manage Subscribers Logging In to ICR Partitions on page 136
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 137
- Monitoring the Configuration of ICR Partitions on page 138

ICR Overview

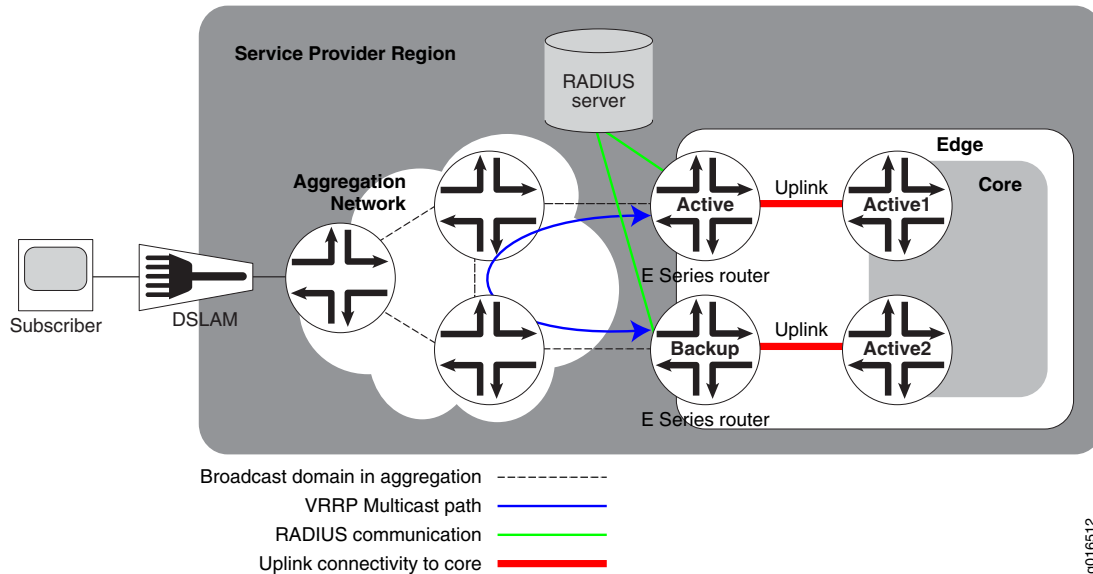
Routers can function as broadband network gateways, which aggregate many subscribers and services simultaneously in order to support intelligent IP services such as voice over IP (VoIP) and video on demand (VoD). If the router fails because of chassis (hardware) failure, subscriber downtime can result.

Interchassis redundancy (ICR) enables you to minimize subscriber downtime when the router or access interface on the edge router fails by re-creating subscriber sessions that were originally terminated on the failed router. It also enables you to track the failure of uplink interfaces. In this way, ICR enables you to completely recover from

router failure. By using extended Virtual Router Redundancy Protocol (VRRP) features, ICR enables you to track the failure of uplink interfaces, elect the master VRRP instance, and detect failure of a VRRP instance.

Figure 7 on page 122 illustrates a suitable network for ICR deployment.

Figure 7: Sample Network for ICR Deployment



Consider subscriber login requests that are received at the digital subscriber line access multiplexer (DSLAM). The request is then forwarded to the aggregation network which forwards the requests to the routers that are part of the edge network. In a conventional network, when the router at the edge network fails, it results in subscriber downtime. By configuring ICR, you can reduce subscriber downtime because the login requests are received by both the *master* (active) and *backup* routers. After authentication from RADIUS, the *master* router accepts the login requests and forwards the same to the uplink interface that is part of the core of the network while the *backup* router rejects the login requests. If the master router fails, login requests are forwarded to the backup router that acts as the master router.

ICR enables load balancing by enabling you to create partitions. An *ICR partition* is a collection of logical subscribers within a single ICR interface. You can manage each ICR partition using a unique VRRP instance.

You can also create ICR clusters. An *ICR cluster* consists of a group of chassis participating in ICR. You can use different E Series routers to configure a heterogeneous ICR cluster. For example, you can use an E120 or E320 router with an ES2 4G LM as a backup for subscribers on an ERX1440 router, or use an ERX1440 router with a GE-HDE LM as a backup for subscribers on an E120 or E320 router. However, you must keep in mind the hardware scaling limitations when you configure an ICR cluster containing both E320 routers and ERX routers.



NOTE: While deploying ICR, service providers must modify the aggregation network to enable subscriber traffic to reach all corresponding interfaces that are configured as part of the redundancy group.

- Related Topics**
- ICR Scaling Considerations on page 124
 - Guidelines for Deploying an ICR Partition in Your Network on page 126
 - Configuring an ICR Partition on page 129

ICR Platform Considerations

ICR is supported on all E Series routers.

For information about modules supported on E120 and E320 routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support ICR.

For information about modules supported on ERX routers:

- See *ERX Module Guide, Table 1, ERX Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support ICR.

Interface Specifiers

The majority of the configuration task examples in this topic collection use the *slot/adapter/port* format to specify an interface. However, the interface specifier format that you use depends on the router that you are using.

For ERX7xx models, ERX14xx models, and ERX310 routers, use the *slot/port* format. For example, the following command specifies a Gigabit Ethernet interface on slot 0, port 1 of an ERX7xx model, ERX14xx model, or ERX310 router.

```
host1(config)#interface gigabitEthernet 0/1
```

For E120 and E320 routers, use the *slot/adapter/port* format, which includes an identifier for the bay in which the I/O adapter (IOA) resides. In the software, adapter 0 identifies the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router). For example, the following command specifies a 10-Gigabit Ethernet interface on slot 5, adapter 0, port 0 of an E320 router.

```
host1(config)#interface tenGigabitEthernet 5/0/0
```

Related Topics ■ Interface Types and Specifiers**ICR Terms**

Table 12 on page 124 defines terms used in this discussion of ICR.

Table 12: ICR Terminology

Term	Description
ICR cluster	Group of E Series routers participating in interchassis redundancy (ICR) deployment.
ICR interface	Physical interface, for example, gigabitEthernet 3/1/3, on an E Series router on which ICR is enabled. The ICR interface is always tied to a unique router.
ICR partition	A logical group of subscriber interfaces within a single ICR interface. For example, the ICR partition can be a group of S-VLANs configured on a single physical interface. You can create multiple partitions on each ICR interface and configure the number of partitions, as well as assign subscribers to the partition. An ICR partition can be configured as master, backup, or dormant.
VRRP	Virtual Router Redundancy Protocol. Use VRRP to prevent loss of network connectivity by configuring backup routers. The backup routers maintain network connectivity when the master router fails. You can configure unique VRRP instances to manage each ICR partition.
VSA	Vendor-specific attributes. VSAs are defined by remote-access server vendors to customize how RADIUS works on their servers. VSAs can be used in combination with RADIUS-defined attributes.

ICR References

For more information about ICR, see the following resources:

- RFC 2338—Virtual Router Redundancy Protocol (April 1998)
- RFC 2787—Definitions of Managed Objects for the Virtual Router (March 2000)
- RFC 2865—Remote Authentication Dial In User Service (RADIUS) (June 2000)
- RFC 2866—RADIUS Accounting (June 2000)

ICR Scaling Considerations

An ICR cluster is a group of routers participating in ICR deployment. When planning an ICR cluster you must ensure that you have adequate resources in the event of a worst-case failure scenario such as a multiple hardware or multiple router failure. For instance, most networks can handle the failure of a single router. However, they may not be able to handle multiple router failures. ICR enables you to choose the

degree of redundancy in your ICR cluster. Depending on the type of network that you have, you can design a 1:1 (minimum) or 1:N (maximum) degree of redundancy in the ICR cluster.




NOTE: Remember to consider parameters such as link bandwidth, QoS, and line module scaling limitations when you plan the deployment of the ICR cluster.

1:1 Subscriber Redundancy in a 4-Node ICR Cluster

Consider a 4-node ICR cluster that consists of four ERX1440 routers, as shown in Figure 8 on page 125. Each of the four routers is capable of supporting 48,000 PPP/PPPoE subscribers. The degree of redundancy that you can achieve in this cluster is 1:1. For every subscriber, you have a backup destination within the cluster. If one router fails, subscriber load is equally distributed to the other three routers. Thus, no single router serves as a dedicated backup. Instead, each router is loaded with 75 percent of its capacity and the remaining 25 percent is kept unused to accommodate subscribers from the failing router. Failure of any one router causes all routers in the cluster to become fully loaded with no spare capacity to accommodate further failures. This is the minimum degree of redundancy in a 4-node ICR cluster.

Figure 8 on page 125 illustrates an example of 1:1 redundancy.

Figure 8: Sample 1:1 Subscriber Redundancy in a 4-Node ICR Cluster



1:3 Subscriber Redundancy in a 4-Node ICR Cluster


You can increase the redundancy in a cluster by configuring more than one backup destination for each subscriber. You can achieve a fully redundant system in which a single router can support all the subscribers when all the other routers fail. To achieve this, you may have to compromise the operating efficiency of the deployed hardware.

Consider the same 4-node ICR cluster but where each router is loaded with 25 percent of its actual capacity as shown in Figure 9 on page 126. If three routers of the

four node cluster fail, the single router has enough spare capability to accommodate the entire subscriber load. This is the maximum degree of redundancy in a 4-node ICR cluster.

Figure 9 on page 126 illustrates an example of 1:3 redundancy.

Figure 9: Sample 1:3 Subscriber Redundancy in a 4-Node ICR Cluster



- Related Topics**
- Guidelines for Deploying an ICR Partition in Your Network on page 126
 - Configuring an ICR Partition on page 129

Guidelines for Deploying an ICR Partition in Your Network

To support 1: N redundancy and to allow your network to scale, you must deploy an ICR partition in your network. An ICR partition is a logical group of subscribers each of which you can manage using a unique VRRP instance. You can configure multiple partitions on a single physical interface to allow your network to distribute subscriber load to multiple backup destinations.

In a stateless model, when a partition fails, all subscribers belonging to the failing partition eventually time out and must log in again to a new active destination based

on the state of the VRRP instance. A partition in the backup state does not accept subscriber login requests. Also, in a stateless model, the router sends early termination requests to clients so that clients do not wait for timeout conditions to occur in order to send requests to log in again. In this way, a stateless model is best suited for PPP/PPPoE subscribers.



NOTE: To avoid routing issues while configuring stateless ICR for PPPoE subscribers, you must configure non-overlapping IP addresses for the clients. This indicates that the client is always assigned a different IP address after failover.

Hardware Requirements for ICR

Before you deploy stateless ICR for PPP/PPPoE subscribers, you must plan and obtain hardware resources to fulfill QoS and bandwidth requirements. You can identify the hardware required based on the type of ICR cluster planned. For instance, if you plan to create a simple ICR cluster that provides 1:1 redundancy, the minimum hardware capability you must plan for is 2. 1:1 redundancy implies that for every subscriber, you have a backup destination within the cluster. Each physical interface configured as part of the ICR cluster must have one backup destination interface on the other router. Also, the router on which you configure the backup partition must have enough bandwidth to accommodate new subscribers after failover.

Network Requirements for ICR

After you plan for and obtain hardware for ICR, you must create the network topology and connections in order to set up a broadcast network between the master and backup interfaces. You can use layer 2 switches and configure them to provide selective broadcast connectivity using VLAN tags. The new broadcast network must allow multicast VRRP traffic between participating ICR interfaces.

Router Configurations for ICR

After setting up the broadcast network, you can configure AAA, QoS, and the interface on the router. You must also configure the uplink interfaces separately on each router that is part of the ICR cluster. Uplink interfaces do not have backups; they behave as if they have been configured on two independent routers.

- Related Topics**
- ICR Scaling Considerations on page 124
 - Configuring an ICR Partition on page 129

Interaction with RADIUS for ICR

You can include an ICR-Partition-Id vendor-specific attribute (VSA) in the following RADIUS messages:

- Access-Request
- Acct-Start
- Acct-Stop

- Interim-Acct (if Acct-Stop messages are specified)
- Partition-Accounting-On
- Partition-Accounting-Off



NOTE: For more information about the ICR partition accounting messages, see the *Configuring RADIUS Attributes* chapter in the *JUNOS Broadband Access Configuration Guide*.

Use the ICR-Partition-Id VSA to determine the ICR partition on which subscribers are logged in. You can configure the same ICR-Partition-Id string for an active ICR partition and its corresponding backup partition.

To configure inclusion of ICR-Partition-Id in RADIUS Access-Request, Acct-Start, and Acct-Stop messages, you can use the ICR-Partition-Id attribute in the **radius include** command. When included in Acct-Stop messages, the attributes are also included in Interim-Acct messages. If you enabled ICR partition accounting and accounting servers are not configured in that partition, it is as if ICR partition accounting is disabled.

In addition to including the ICR-Partition-Id VSA in RADIUS Access-Request, Acct-Start, Acct-Stop, and Interim-Acct messages, the router also sends the Partition-Accounting-On and Partition-Accounting-Off messages:

Both Partition-Accounting messages include the ICR-Partition-Id VSA. Also, both these messages are sent to the RADIUS accounting server configured on the virtual router where the ICR partition is configured or the virtual router on which the ICR control interface is set up.

You can optionally configure duplicate or broadcast AAA accounting on a virtual router, which sends the accounting information to additional virtual router simultaneously, so that the Partition-Accounting-On and Partition-Accounting-Off messages can also be sent to the duplicate and broadcast virtual routers.

ICR Partition Accounting Overview

To enable or disable sending of the ICR Partition-Accounting-On or Partition-Accounting-Off messages to the RADIUS servers, you can now use the **radius icr-partition-accounting** command. If you enabled ICR partition accounting on a virtual router, subscribers are allowed to log in to that partition only after the response for the corresponding Partition-Accounting-On message is received from the RADIUS server. As a result, the time that it takes to receive a response from the RADIUS server might increase the time it takes for subscribers to log in. Also, when an ICR partition transitions from the master to backup state, the Partition-Accounting-Off message is sent to the RADIUS server only after the router receives the response for the Acct-Stop messages for all the subscribers logged out in that partition.

The transition of the ICR partition states from master to backup and backup to master can occur because of chassis failure, an administrative switchover, or an interface or line module reset action. The following scenarios describe how ICR partition accounting messages are processed and subscriber logging is handled:

- In the event of a complete chassis failure, RADIUS does not interact with the failing B-RAS application on the router. In such a scenario, when the new master partition takes over, the Partition-Accounting-On message is sent from the new master. After the response for the Partition-Accounting-On message is received from the new master partition, subscribers are allowed to log in to the master. When you remove certain VLAN or S-VLAN IDs from an ICR partition, the corresponding subscribers in that partition are removed and forced to log out from the chassis. This action causes the Acct-Stop messages to be sent to RADIUS.
- If ICR partition accounting is enabled and an administrative switchover forces subscribers in a particular ICR partition to logged out, the Partition-Accounting-Off message is sent from the failing B-RAS application on the router only after Acct-Stop responses are received for all the logged out subscribers.
- If ICR partition accounting is enabled, and the interface or the line module that is configured with the ICR partition resets, the Partition-Accounting-Off message is sent from the failing B-RAS application on the router after Acct-Stop responses are received for all subscribers in that partition.

- Related Topics**
- Using RADIUS to Manage Subscribers Logging In to ICR Partitions on page 136
 - radius include
 - radius icr-partition-accounting
 - show radius icr-partition-accounting
 - *Configuring RADIUS Attributes in the JUNOS Broadband Access Configuration Guide*

Configuring an ICR Partition

When you configure an ICR partition, you configure the interface on which the ICR partition resides and create a unique VRRP instance to manage the partition.

To configure an ICR partition:

1. Configure the interface.

See “Configuring the Interface on Which the ICR Partition Resides” on page 130.
2. Create a unique VRRP instance to manage the ICR partition.

See “Configuring VRRP Instances to Match ICR Requirements” on page 131.
3. Create and assign a name to the ICR partition.

See “Naming ICR Partitions” on page 131.
4. (Optional) Group the subscribers.

See “Grouping ICR Subscribers Based on S-VLAN IDs” on page 132 and “Grouping ICR Subscribers Based on VLAN IDs” on page 133.



NOTE: Grouping subscribers based on S-VLAN IDs is the default grouping option for ICR partitions. If you do not explicitly specify the grouping option, subscribers are grouped based on S-VLAN IDs.

5. (Optional) Configure RADIUS.

See “Using RADIUS to Manage Subscribers Logging In to ICR Partitions” on page 136

Related Topics

- ICR Overview on page 121
- Monitoring the Configuration of ICR Partitions on page 138
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 137
- Monitoring the Status of ICR Partition Accounting

Configuring the Interface on Which the ICR Partition Resides

You can create multiple ICR partitions on an interface. For information on the number of ICR partitions that you can create, see *JUNOS Release Notes, Appendix A, System Maximums*. However, you cannot create a single ICR partition that spans multiple interfaces.

To configure the interface on which the ICR partition resides:

1. Specify a FastEthernet, GigabitEthernet, or 10-GigabitEthernet interface.

```
host1(config)#interface gigabitEthernet 3/5/0
host1(config-if)#
```

2. Specify VLAN as the encapsulation method to create the VLAN major interface.

```
host1(config-if)#encapsulation vlan
```

3. Create a VLAN subinterface by adding a subinterface number to the interface identification number.

```
host1(config-if)#interface gigabitEthernet 3/5/0.10
```

4. Assign a VLAN ID for the subinterface. The router configures the subinterface whether or not the subinterface is part of the ICR partition. Use the **icr-control-interface** keyword to specify that an ICR partition can be configured on the the subinterface.

```
host1(config-if)#vlan id 10 1 icr-control-interface
```

5. Assign an IP address to the VLAN subinterface.

```
host1(config-if)#ip address 3.5.1.1/24
```


- Related Topics**
- Configuring VRRP Instances to Match ICR Requirements on page 131
 - Monitoring the Configuration of an ICR Partition Attached to an Interface on page 137

Configuring VRRP Instances to Match ICR Requirements

Each ICR partition is managed by a unique VRRP instance. You can configure an ICR partition as the *master* partition using the **ip vrrp priority** command. VRRP is also used to detect failure in the uplink interfaces.

To configure the VRRP instance to match ICR requirements:

1. Create a VRRP instance by specifying the identification number, and associate an IP address with the identification number.

```
host1(config-if)#ip vrrp 1 virtual-address 3.5.1.10
```

2. Specify the priority of the router. Assign the higher priority to the master ICR partition and a lower priority to the backup ICR partition.

```
host1(config-if)#ip vrrp priority 200
```

3. Enable the router to learn the VRRP advertisement interval.

```
host1(config-if)#ip vrrp 1 timers-learn
```

4. Enable the VRRP instance.

```
host1(config-if)#ip vrrp 1 enable
```

5. (Optional) Configure additional VRRP instances by completing Steps 1 through 4, using unique numbering.

Each ICR partition is managed by a unique VRRP instance. Configure additional VRRP instances only if you plan to create additional ICR partitions.

- Related Topics**
- *Chapter 4, Configuring VRRP*
 - ip vrrp
 - ip vrrp enable
 - ip vrrp priority
 - ip vrrp timers-learn
 - ip vrrp virtual-address

Naming ICR Partitions

After you have configured the interface on which the ICR partition resides and the unique VRRP instance that manages the ICR partition, you must create the ICR

partition. You can use the keywords *master* or *backup* to identify the type of ICR partition created.

To create and name ICR partitions:

1. Create an ICR partition by specifying a unique name for the partition. For easy identification, you can include the keywords *master* or *Backup* in the name of the partition.

```
host1(config-if)#ip vrrp 1 icr-partition part1Master
```

2. (Optional) Create additional ICR partitions by repeating Step 1, using unique names or numbering.



NOTE: You can manage each ICR partition using a unique VRRP instance. Before you create additional ICR partitions, you must create corresponding VRRP instances.

```
host1(config-if)#ip vrrp 2 icr-partition part1Backup
```

```
host1(config-if)#ip vrrp 3 icr-partition part1Dormant
```

For information on the number of ICR partitions that you can create per line module or chassis, see *JUNOS Release Notes, Appendix A, System Maximums*.

Related Topics

- ICR Overview on page 121
- ip vrrp icr-partition
- Monitoring the Configuration of ICR Partitions on page 138

Grouping ICR Subscribers Based on S-VLAN IDs

You can group ICR subscribers based on S-VLAN IDs. When you configure an S-VLAN list or S-VLAN range or an S-VLAN and VLAN subinterface pair, you can include any or all of the following keywords:

- Use the **control-interface** keyword to control traffic on the range of subinterfaces. If the subinterfaces are part of the backup partition, the router changes the state of all the subinterfaces to AdminDown, blocking all traffic to the subinterfaces.
- Use the **use-default-mac** keyword to enable the subinterfaces to use the default MAC address instead of the VRRP MAC address. To enable the subscriber subinterface to use the correct MAC address, configure the subinterface after configuring the ICR partition.
- Use the **advertise-mac** keyword to enable the subinterfaces to transmit gratuitous ARP (GARP) advertisements when the ICR partition moves from the backup state to the master state.

To group ICR subscribers based on S-VLAN IDs:

1. Specify **svlan** as the grouping type.

```
host1(config-if)#ip vrrp 1 icr-partition group svlan
```

The default grouping option is S-VLAN. If you do not explicitly specify the grouping option, the subscribers are grouped based on S-VLAN.

2. Add S-VLAN subinterfaces to the ICR partition by doing either of the following:
 - Specify the S-VLAN IDs individually by using the **svlan-list** keyword. In the following example, you add individual S-VLAN subinterfaces by specifying each S-VLAN ID.

```
host1(config-if)#ip vrrp 1 icr-partition svlan-list 100 102 105 108 114 125
control-interface use-default-mac advertise-mac
```

- Specify the starting ID and ending ID of the range of S-VLAN subinterfaces. In the following example, you specify the first and the last ID of the range because the IDs are in sequential order.

```
host1(config-if)#ip vrrp 1 icr-partition svlan-range 100 110 control-interface
use-default-mac advertise-mac
```

3. (Optional) Add an S-VLAN and VLAN subinterface pair to the ICR partition.

```
host1(config-if)#ip vrrp 1 icr-partition svlan-list-explicit 120 1 120 2
control-interface use-default-mac advertise-mac
```

4. (Optional) Configure additional S-VLAN subinterfaces by completing Steps 2 and 3 using unique numbering.

Related Topics

- Grouping ICR Subscribers Based on VLAN IDs on page 133
- ip vrrp icr-partition group
- ip vrrp icr-partition svlan-list
- ip vrrp icr-partition svlan-list explicit
- ip vrrp icr-partition svlan-range
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 137

Grouping ICR Subscribers Based on VLAN IDs

You can configure ICR subscribers based on VLAN IDs. When you configure a VLAN list or VLAN range, you can include any or all of the following keywords:

- Use the **control-interface** keyword to control traffic on the range of subinterfaces. If the subinterfaces are part of the backup partition, the router changes the state of all the subinterfaces to AdminDown, thus blocking all traffic to the subinterfaces.
- Use the **use-default-mac** keyword to enable the subinterfaces to use the default MAC address instead of the VRRP MAC address. To enable the subscriber subinterface to use the correct MAC address, configure the subinterface after configuring the ICR partition.

- Use the **advertise-mac** keyword to enable the subinterfaces to transmit gratuitous ARP (GARP) advertisements when the ICR partition moves from the backup state to the master state.

To group ICR subscribers based on VLAN IDs:

1. Specify VLAN as the grouping type.

```
host1(config-if)#ip vrrp 1 icr-partition group vlan
```

The default grouping option is S-VLAN. If you do not explicitly specify the grouping option, the subscribers are grouped based on S-VLAN.

2. Add VLAN subinterfaces to the ICR partition by doing either of the following:

- Specify the VLAN IDs individually by using the **vlan-list** keyword to add a group of random VLAN IDs. In the following example, you add VLAN subinterfaces by specifying each VLAN ID individually because the IDs are in random order.

```
host1(config-if)#ip vrrp 1 icr-partition vlan-list 10 21 62 control-interface
use-default-mac advertise-mac
```

- Specify the starting ID and ending ID of the range of VLAN subinterfaces. In the following example, you specify the first and the last ID of the range because the IDs are in sequential order.

```
host1(config-if)#ip vrrp 1 icr-partition vlan-range 10 40 control-interface
use-default-mac advertise-mac
```

3. (Optional) Configure additional VLAN subinterfaces by completing Step 2 using unique numbering.

Related Topics

- Grouping ICR Subscribers Based on S-VLAN IDs on page 132
- ip vrrp icr-partition group
- ip vrrp icr-partition vlan-list
- ip vrrp icr-partition vlan-range
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 137

Example: Configuring ICR Partitions That Group Subscribers by S-VLAN ID

The following example show how to configure a *master* ICR partition on an ERX1440 router. In this example, you first configure the interface on which the ICR partition resides. You can then create a new VRRP instance to manage the ICR partition. The value you assign to the **priority** keyword determines the state of the ICR partition. The VRRP instance that has the higher priority manages the master ICR partition. All other VRRP instances manage the backup and dormant partitions.

1. Configure the interface on which the ICR partition resides.

```
host1 (config)#interface gigabitEthernet 3/5
```

```

host1 (config-if)#encapsulation vlan
host1 (config-if)#interface gigabitEthernet 3/5.10
host1 (config-if)#svlan id 10 1 icr-control-interface
host1 (config-if)#ip address 3.5.1.1/24

```

2. Configure the VRRP instance based on the ICR partition requirements.

```

host1 (config-if)#ip vrrp 1 virtual-address 3.5.1.10
host1 (config-if)#ip vrrp 1 priority 200
host1 (config-if)#ip vrrp 1 timers-learn
host1 (config-if)#ip vrrp 1 enable

```

3. Create and identify the ICR partition.

```

host1 (config-if)#ip vrrp 1 icr-partition part1Master

```

4. Group subscribers based on S-VLAN IDs.

```

host1 (config-if)#ip vrrp 1 icr-partition group svlan
host1 (config-if)#ip vrrp 1 icr-partition svlan-range 100 110 control-subinterface
host1 (config-if)#ip vrrp 1 icr-partition svlan-range 111 119 advertise-mac
use-default-mac
host1 (config-if)#ip vrrp 1 icr-partition svlan-list-explicit 120 1 120 2
advertise-mac control-subinterface use-default-mac
host1 (config-if)#exit

```

The following example shows how to configure a *backup* ICR partition on an E320 router. Configure the interface on which the ICR partition resides and then create a new VRRP instance that manages the backup ICR partition. The value you assign to the **priority** keyword determines the state of the ICR partition. In the case of a backup ICR partition, specify a value lower than the priority of the master ICR partition.

1. Configure the interface on which the ICR partition resides.

```

host2 (config)#interface gigabitEthernet 11/1/0
host2 (config-if)#encapsulation vlan
host2 (config-if)#interface gigabitEthernet 11/1/0.10
host2 (config-if)#svlan id 10 1 icr-control-interface
host2 (config-if)#ip address 3.5.1.2/24

```

2. Configure the VRRP instance based on the ICR partition requirements.

```

host2 (config-if)#ip vrrp 1 virtual-address 3.5.1.10
host2 (config-if)#ip vrrp 1 priority 100
host2 (config-if)#ip vrrp 1 timers-learn
host2 (config-if)#ip vrrp 1 enable

```

3. Create and identify the ICR partition.

```

host2 (config-if)#ip vrrp 1 icr-partition part1Backup

```

4. Group subscribers based on S-VLAN IDs.

```

host2 (config-if)#ip vrrp 1 icr-partition group svlan
host2 (config-if)#ip vrrp 1 icr-partition svlan-range 100 110 control-subinterface

```

```

host2 (config-if)#ip vrrp 1 icr-partition svlan-range 111 119 advertise-mac
use-default-mac
host2 (config-if)#ip vrrp 1 icr-partition svlan-list-explicit 120 1 120 2
advertise-mac control-subinterface use-default-mac
host2 (config-if)#exit

```

Grouping subscribers based on S-VLAN IDs is the default grouping method for ICR partitions. You can also explicitly choose S-VLAN as the grouping option as shown in this example. To add a group of random S-VLAN IDs, use the **svlan-list** command.

To group subscribers by VLAN IDs, use the **vlan** keyword instead of the **svlan** keyword. To add a group of random VLAN IDs, use the **vlan-list** command.



NOTE: While grouping subscribers based on VLAN IDs, you can use corresponding VLAN grouping commands. However, the **svlan-list-explicit** command does not have any corresponding VLAN command.

-
- Related Topics**
- ICR Overview on page 121
 - Guidelines for Deploying an ICR Partition in Your Network on page 126
 - ICR Scaling Considerations on page 124

Using RADIUS to Manage Subscribers Logging In to ICR Partitions

To configure RADIUS to manage subscribers logging in to ICR partitions on the router, perform the following tasks:

- Configure inclusion of the ICR-Partition-ID VSA in RADIUS messages.

```
host1(config)#radius-include icr-partition-id acct-start enable
```

Issuing this command includes the ICR-Partition-ID VSA in Acct-Start messages. To include the ICR-Partition-ID VSA in other accounting and access messages, see the *Configuring RADIUS Attributes* chapter in the *JUNOS Broadband Access Configuration Guide*.

- Enable or disable sending of the ICR Partition-Accounting-On or Partition-Accounting-Off messages to the RADIUS servers.

```
host1(config)#radius icr-partition-accounting enable
```

For more information on enabling or disabling sending of partition accounting messages to RADIUS servers configured on a virtual router, see the *Configuring RADIUS Attributes* chapter in the *JUNOS Broadband Access Configuration Guide*.

-
- Related Topics**
- radius include
 - radius icr-partition-accounting
 - show radius icr-partition-accounting

- Interaction with RADIUS for ICR on page 127
- Configuring an ICR Partition on page 129

Monitoring the Configuration of an ICR Partition Attached to an Interface

Purpose Display information about the ICR partition configured on an interface.

Action `host1#show icr-partition fastEthernet 3/5/0.1 1`
 ICR Partition ID: part1A
 ICR Partition State: Master
 ICR Partition Grouping Criterion: SVLAN

SVLAN	VLAN	control-interface	vrrp-mac	advertise-mac
100	Any	enabled	disabled	enabled
101	Any	enabled	disabled	disabled
102	Any	enabled	disabled	disabled
103	Any	enabled	disabled	disabled
104	Any	enabled	disabled	disabled
105	Any	enabled	disabled	disabled
106	Any	enabled	disabled	disabled
107	Any	enabled	disabled	disabled
108	Any	enabled	disabled	disabled
109	Any	enabled	disabled	disabled

ICR Partition has 10 group members.

Meaning Table 13 on page 137 lists the `show icr-partition` command output fields.

Table 13: show icr-partition Output Fields

Field Name	Field Description
ICR Partition ID	Identifier for the ICR partition.
ICR Partition State	State of the ICR partition: <ul style="list-style-type: none"> ■ Master—ICR partition that accepts subscriber login requests. ■ Backup—ICR partition that does not accept subscriber login requests. ■ Dormant—When the IP address or virtual router is forcibly deleted, or if the lower interface is not available, the ICR partition moves to the Dormant state. The dormant ICR partition does not accept subscriber login requests. <p>NOTE: The state of the ICR partition depends on the associated VRRP instance.</p>
ICR Partition Grouping Criterion	Grouping option for the subscribers. Possible options: S-VLAN and VLAN. The default grouping option is S-VLAN.
SVLAN	S-VLAN identifier for the interface.

Table 13: show icr-partition Output Fields (continued)

Field Name	Field Description
VLAN	VLAN identifier for the interface. Any indicates that the VLAN ID is a wildcard and you can specify any configured VLAN ID with the associated S-VLAN ID.
control-interface	Controls traffic on the interface. Possible states: enabled or disabled. If the status is enabled and the interface is part of the backup partition, the router changes the state of the interface to Admindown and blocks all traffic to the interface. If the status is disabled, the router does not control traffic on the interface.
vrrp-mac	Configures the interface to use the default MAC address instead of the VRRP MAC address. Possible states: enabled or disabled. If the status is enabled, the interface uses the default MAC address; otherwise, the interface uses the VRRP MAC address.
advertise-mac	Enables the interface to transmit GARP advertisements when the partition moves from backup state to master state. Possible states: enabled or disabled. If the status is enabled, the interface transmits GARP advertisements; otherwise, the interface does not transmit GARP advertisements.

- Related Topics**
- Configuring the Interface on Which the ICR Partition Resides on page 130
 - show icr-partition

Monitoring the Configuration of ICR Partitions

Purpose Display information about ICR partitions and their status.

Action To display information about all ICR partitions:

```

host1#show icr-partitions
Interface-Location Vrrp-Id   State      Partition-ID
-----
3/5/0.2           20      *Backup    part20A
3/5/0.1           10      Master     part10A
2/1/0.1           1       Backup     part1Backup
2/5/0.2           2       Backup     part2Backup
3/1/0.1           4       Dormant    part4
-----
Total ICR Partitions: 5

```

To display information based on the state of the ICR partition:

```

host1#show icr-partitions Master
Interface-Location Vrrp-Id   State      Partition-ID

```



```

-----
3/5/0.1          10      Master      part10A
-----
Total ICR Partitions in Master state: 1

```

To display a summary of the ICR partitions configured:

```

host1#show icr-partitions summary
Dormant ICR Partitions: 1
Backup ICR Partitions: 3
Master ICR Partitions: 1
Total ICR Partitions: 5

```

You can also display information about configured ICR partitions using a filter as an alternative to specifying the **state** keyword. For instance, to display information about the backup and dormant ICR partitions only, you can use the **exclude Master** keywords, as shown in the following example:

```

host1#show icr-partitions | exclude Master
Interface-Location Vrrp-Id   State      Partition-ID
-----
3/5/0.2           20      *Backup    part20A
2/1/0.1           1       Backup     part1Backup
2/5/0.2           2       Backup     part2Backup
3/1/0.1           4       Dormant     part4
-----
Total ICR Partitions: 5

```

Meaning Table 14 on page 139 lists the **show icr-partitions** command output fields.

Table 14: show icr-partitions Output Fields

Field Name	Field Description
Interface-Location	Interface Identifier or location identifier of the ICR partition.
Vrrp-Id	VRRP identifier of the VRRP instance associated with the ICR partition.
State	<p>State of the ICR partition:</p> <ul style="list-style-type: none"> ■ Master—ICR partition that accepts subscriber login requests. ■ Backup—ICR partition that does not accept subscriber login requests. ■ Dormant—When the IP address or virtual router is forcibly deleted, or if the lower interface is not available, the ICR partition moves to the Dormant state. The dormant ICR partition does not accept subscriber login requests. <p>NOTE: The state of the ICR partition depends on the associated VRRP instance. When the state of the VRRP instance changes, the state of the ICR partition also changes. A '*' associated with an ICR partition indicates that the partition is in transition.</p>

Table 14: show icr-partitions Output Fields *(continued)*

Field Name	Field Description
Partition-ID	Identifier for the ICR partition.
Dormant ICR Partitions	Number of dormant ICR partitions configured on the router.
Backup ICR Partitions	Number of backup ICR partitions configured on the router.
Master ICR Partitions	Number of master ICR partitions configured on the router.
Total ICR Partitions	Total number of ICR partitions configured on the router.

- Related Topics**
- [Configuring the Interface on Which the ICR Partition Resides](#) on page 130
 - [show icr-partitions](#)

Part 2

Index

- Index on page 143

Index

A

Access-Request messages	
ICR Partition ID VSA.....	128
Acct-Start messages	
ICR Partition ID VSA.....	128
Acct-Stop messages	
ICR Partition ID VSA.....	128
assembly numbers, displaying for hardware.....	19
assembly revisions, displaying for hardware.....	19
automatic switchover.....	9

B

backup router.....	106
defined.....	102
election process and.....	106
VRRP.....	101
bandwidth	
optimizing.....	7
baseline commands	
baseline ip vrrp.....	113

C

clear commands	
clear redundancy history.....	46
conventions	
notice icons.....	xix
text and syntax.....	xx
customer support.....	xxi
contacting JTAC.....	xxi

D

destination address (DA), VRRP.....	103
disable-switch-on-error command.....	16
documentation set	
comments on.....	xxi

F

failover. <i>See</i> switchover	
---------------------------------	--

H

hardware	
monitoring information.....	19
high availability	
activating.....	40
deactivating.....	41
IP interface priority.....	41, 42
monitoring.....	43
overview.....	25

I

icr cluster.....	124
ICR commands	
interface	130
icr interface.....	124
ICR Options	
icr-control-interface.....	130
priority command.....	131
timers-learn command.....	131
ICR Partition	
configuring	129, 130, 131, 132, 133
configuring, naming.....	131
radius.....	136
ICR partition accounting	
and dependence on Acct-Stop messages.....	128
configuring.....	127
disabling and enabling messages	
sent to the RADIUS server.....	127
overview.....	127
processing in different scenarios	
administrative switchover.....	128
complete chassis failure.....	128
line module or interface failure.....	128
transition of ICR partition states.....	128
ICR Partition commands	
naming	132
ICR partition ID VSA	
including in access and accounting	
messages.....	128

ICR Partition ID VSA	
transmitting to the virtual router	
where ICR control interface is	
configured.....	128
where ICR partition is configured.....	128
ICR Partition Options	
advertise-mac.....	133, 134
control-subinterface	133, 134
group option.....	132, 134
use-default-mac	133, 134
ICR Partitions	
configuration example	
backup ICR partition, S-VLAN based	
grouping.....	135
master ICR partition, S-VLAN based	
grouping.....	134
ICR RADIUS commands	
inclusion of icr-partition-id.....	136
radius icr-partition-accounting.....	136
icr-partition.....	124
in-service software upgrade. <i>See</i> unified ISSU	
Interchassis Redundancy	
heterogenous icr clusters.....	122
icr clusters.....	122
icr partition.....	122
Interim-Acct messages	
ICR Partition ID VSA.....	128
IP addresses	
IP address owner, VRRP.....	102
primary, VRRP.....	102
VRRP.....	106, 108
ip commands.....	108
ip initial-sequence-preference.....	42
ip vrrp.....	108
ip vrrp accept-data.....	108
ip vrrp advertise-interval.....	108
ip vrrp authentication-key.....	108
ip vrrp authentication-type.....	108
ip vrrp enable.....	108
ip vrrp preempt.....	108
ip vrrp priority.....	108
ip vrrp track.....	112
ip vrrp virtual-address.....	108
<i>See also</i> vrrp commands	
ip pim commands	
ip pm dr-priority.....	84
ipv6 commands	
ipv6 initial-sequence-preference.....	42
ISSU. <i>See</i> unified ISSU	

L

LEDs	
monitoring status.....	19

line module redundancy	
configuring.....	7, 10
E120 and E320 Broadband Services Routers.....	7
IOA behavior.....	7
ERX7xx models and ERX14xx models.....	7
managing.....	10
monitoring.....	19

M

manuals	
comments on.....	xxi
master router.....	102
memory (hardware), displaying.....	19

N

notice icons.....	xix
-------------------	-----

O

optimizing bandwidth.....	7
overload advertise-high-metric issu command.....	81, 83

P

physical slots	
rebooting.....	15
platform considerations	
high availability.....	26

R

RADIUS.....	124
redundancy	
line module. <i>See</i> line module redundancy	
SRP module. <i>See</i> SRP module redundancy	
redundancy commands	
redundancy force-switchover.....	10, 16
redundancy lockout.....	10
redundancy revert.....	11
redundancy revertive.....	10
reversion	
after switchover.....	7
revisions, displaying assembly.....	19

S

serial numbers, displaying for hardware.....	19
Service Availability	
Features.....	5
ICR.....	6
Module redundancy.....	5
Stateful SRP Switchover.....	5
Unified ISSU.....	5
VRRP.....	6

show environment command.....	19
show hardware command.....	20
show icr commands	
show icr-partition.....	137
show icr-partitions.....	138
state.....	138
summary.....	139
show ip commands	
show ip vrrp.....	113
show ip vrrp brief.....	113
show ip vrrp neighbors.....	113
show ip vrrp statistics.....	113
show ip vrrp statistics global.....	113
show ip commands:show	
ip interface.....	43
show issu command.....	97
show redundancy commands	
show redundancy.....	19, 44, 46
show redundancy clients.....	46
show redundancy switchover-history.....	46
show version command.....	22
software	
upgrading.....	19
SRP module redundancy.....	11
managing.....	16
monitoring.....	19
SRP modules	
installing a redundant module.....	14
reset button.....	11
srp switch command.....	18
stateful SRP switchover.	27
<i>See also</i> high availability	
status LEDs, monitoring.....	19
support, technical <i>See</i> technical support	
switchover.....	7
synchronize command.....	16, 17

T

technical support	
contacting JTAC.....	xxi
text and syntax conventions.....	xx

U

unified ISSU (in-service software upgrade).....	55
AAA support.....	74
application support.....	65
application-specific behavior.....	73
ATM support.....	74
ATM port data rate.....	74
ILMI sessions.....	74
OAM CC effects.....	74
OAM VC integrity.....	74
VC and VP statistics.....	74

DHCP support.....	75
common component.....	75
external server.....	75
packet capture.....	75
relay and relay proxy.....	75
DHCP support:relay and relay proxy.....	66
DoS protection support.....	75
Ethernet support.....	76
ARP entries.....	76
LAG.....	76
port data rate.....	76
VLAN statistics.....	76
FTP support.....	77
halting during initialization.....	95
halting during upgrade.....	95
initialization phase.....	58
application data on standby SRP	
module.....	58
line module arming.....	58
SNMP traps.....	58
IS-IS support.....	79
graceful restart.....	79
high link cost.....	79
L2TP support.....	81
layer 3 protocol traffic forwarding.....	86
monitoring.....	97
OSPF support.....	82
dead interval.....	82
graceful restart.....	82
high link cost.....	82
overview.....	55
phases	
initialization.....	58
overview of.....	58
service restoration.....	58
upgrade.....	58
PIM support.....	84
platform.....	57
procedure for upgrade.....	91
references.....	58
requirements	
hardware.....	89
software.....	89
traffic forwarding.....	89
verification in upgrade phase.....	58
restoring original router state.....	95
router behavior.....	55
service restoration phase.....	58
SONET/SDH support.....	85
subscriber support.....	84
logins.....	84
statistics.....	84
support, application.....	65
T3 support.....	85
TACACS+ support.....	85
terms.....	57

timer settings for routing protocol timers.....	88
upgrade phase.....	58
exceptions.....	58
line module control plane.....	58
line module forwarding plane upgrade.....	58
process steps.....	58
setup.....	58
SRP module switchover.....	58
verification requirements.....	58
upgrade procedure.....	91
unified ISSU (in-service software upgrade):DHCP support	
local server.....	66
upgrading software.....	19

V

virtual MAC address.....	101
virtual router ID (VRID). <i>See</i> VRID	
Virtual Router Redundancy Protocol (VRRP). <i>See</i> VRRP	
VRID (virtual router ID)	
configuration.....	108
creating.....	108
router election rules.....	106
VRRP.....	124
VRRP (Virtual Router Redundancy Protocol)	
advertisement interval.....	108
advertisement messages.....	106
authentication key.....	108
authentication type.....	108
backup router.....	101, 106
configuration examples.....	103
configuring.....	108
how it works.....	103
implementation.....	106
MAC address.....	101
master router.....	102
monitoring.....	113
overview.....	101
preemption.....	106, 108
router election rules.....	106
router priority.....	108
VLAN support.....	101
VRRP router defined	102
vrrp commands	
ip vrrp.....	108
ip vrrp accept-data.....	108
ip vrrp advertise-interval.....	108
ip vrrp authentication-key.....	108
ip vrrp authentication-type.....	108
ip vrrp enable.....	108
ip vrrp preempt.....	108
ip vrrp priority.....	108
ip vrrp track.....	112
ip vrrp virtual-address.....	108

VRRP commands	
icr.....	131

W

warm restart	
IP interface priority during.....	41, 42