

JunosE™ Software for E Series™ Broadband Services Routers

Service Availability Configuration Guide

Release

11.3.x



Published: 2010-10-08

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

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

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

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 11.3.x

Copyright © 2010, 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 2010—FRS JunosE 11.3.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 OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

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

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

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

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

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

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

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the

Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

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

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

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

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

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

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

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

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

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

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

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Abbreviated Table of Contents

	About the Documentation	xix
Part 1	Chapters	
Chapter 1	Service Availability	3
Chapter 2	Managing Module Redundancy	9
Chapter 3	Managing Stateful SRP Switchover	35
Chapter 4	Managing Stateful Line Module Switchover	67
Chapter 5	Configuring a Unified In-Service Software Upgrade	101
Chapter 6	Configuring VRRP	149
Chapter 7	Managing Interchassis Redundancy	169
Part 2	Index	
	Index	189

Table of Contents

	About the Documentation	xix
	E Series and JunosE Documentation and Release Notes	xix
	Audience	xix
	E Series and JunosE 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
	Module Redundancy	5
	Stateful SRP Switchover	5
	Stateful Line Module Switchover	5
	Unified ISSU	6
	VRRP	6
	Interchassis Redundancy	6
Chapter 2	Managing Module Redundancy	9
	Line Module Redundancy Overview	9
	Line Module Redundancy Requirements	10
	ERX7xx Models and ERX14xx Models	10
	E120 and E320 Routers	10
	IOA Behavior When the Router Reboots	11
	Line Module Behavior When Disabling or Enabling IOAs	11
	Understanding Automatic Switchover	12
	Limitations of Automatic Switchover	12
	Understanding Reversion After Switchover	12
	Configuring Line Module Redundancy	13
	Managing Line Module Redundancy	13
	Example: Forcing the Router to Switch from Primary Line Module to Spare Line Module	14
	Interoperation of Redundancy and Stateful Switchover for Line Modules	15
	Understanding SRP Module Redundancy	16
	Understanding Configuration of SRP Modules for Redundancy	19
	Installing a Redundant SRP Module	20

	Managing SRP Module Redundancy	21
	Switching to the Redundant SRP Module	22
	Determination of Redundancy Status for Line Modules and SRP Modules Using Status LEDs	22
	Monitoring Redundancy in Installed Hardware	23
	Monitoring Redundancy in Line Module and SRP Modules	27
	Monitoring Redundancy Status on E320 Router	30
Chapter 3	Managing Stateful SRP Switchover	35
	Stateful SRP Switchover Overview	35
	Stateful SRP Switchover Platform Considerations	36
	Module Requirements	36
	Stateful SRP Switchover Redundancy Modes	37
	File System Synchronization Mode	37
	High Availability Mode	37
	Stateful SRP Switchover States	38
	Disabled State	39
	Initializing State	40
	Active State	40
	Pending State	41
	Application Support for Stateful SRP Switchover	42
	Application Support	42
	Guidelines for Activating High Availability	51
	Activating High Availability	52
	Guidelines for Deactivating High Availability	53
	Deactivating High Availability	53
	Guidelines for Setting the IP Interface Priority	54
	Setting the IP Interface Priority	54
	Guidelines for Upgrading Software	55
	Monitoring the Redundancy Status	56
	Monitoring the Redundancy Status of Applications	59
	Monitoring the Redundancy History	61
	Monitoring the Redundancy Status of Line Modules	62
	Monitoring the Redundancy Status of SRP Modules	63
	Monitoring the Redundancy Switchover History	65
	Clearing the Redundancy History	65
Chapter 4	Managing Stateful Line Module Switchover	67
	Stateful Line Module Switchover Overview	68
	Benefits of Stateful Line Module Switchover	68
	1:1 Redundancy Model	69
	Seamless Preservation of Subscriber Sessions	69
	Stateful Line Module Switchover Platform Considerations	70
	Guidelines for Configuring Stateful Line Module Switchover	70
	System Operations When Stateful Line Module Switchover Is Enabled	74
	Stateful Line Module Configuration Scenarios	75
	High Availability Configured and Enabled on the Line Module	75
	High Availability Configured and Disabled on the Line Module	76
	High Availability Configured and the Switchover State Is Active or Disabled	76

Rebooting of the System When Line Module High Availability Is Configured	76
Stateful SRP Switchover	76
Line Module Redundancy	77
Unified ISSU	77
Replacement of Line Modules When Stateful Line Module Switchover Is Enabled	77
Reloading the Primary Line Module in Response to Failures	77
Reloading the Secondary Line Module in Response to Failures	78
Disabling the Primary and Secondary Line Module Slots	78
Reloading the Router When Line Modules Enabled for HA Are Installed . . .	78
Removing IOAs Without Powering Down from Line Modules	78
Cold and Warm Switchovers of Line Modules In a High Availability Pair . . .	78
Application Support for Stateful Line Module Switchover	79
Policy Management	79
QoS	80
Connection Manager and Queue Manager	80
PPP	80
L2TP	81
Forwarding Controller	81
Mirroring Subsystem	82
Unified ISSU	82
ICCP	83
Stateful Line Module Switchover Modes	83
Stateless Switchover Mode	83
High Availability Mode	83
Stateful Line Module Switchover States	84
Disabled State	85
Initializing State	86
Active State	86
Guidelines for Activating High Availability	87
Activating High Availability	88
Guidelines for Deactivating High Availability	89
Deactivating High Availability	89
Switching Over from a Primary Line Module to Secondary Line Module	90
Log Messages Generated for Stateful LM Switchover	91
Log Messages Displayed During the Transition from Disabled State to Active State	91
Log Messages Displayed During the Transition from Active State to Pending or Disabled State	91
Log Messages Displayed During the Transition from Pending or Disabled State to Active State	92
Log Messages Displayed During the Transition from Active or Pending State to Disabled State	92
Log Messages Displayed for Stateful SRP and Line Module Switchover When HA Is Enabled	92
Log Messages Displayed for Stateful SRP and Line Module Switchover When HA Is Disabled	93

	Preservation of Statistics During Stateful Line Module Switchover	93
	PPP Accounting Statistics	93
	Policy Statistics	94
	Performance Impact and Scalability Considerations	94
	Use of Status LEDs to Monitor the High Availability States of Line Modules	95
	Monitoring the Redundancy Status of Line Modules in a Specific Slot	95
	Monitoring the Redundancy History of Line Modules in a Specific Slot	97
Chapter 5	Configuring a Unified In-Service Software Upgrade	101
	Unified ISSU Overview	102
	Router Behavior During a Unified In-Service Software Upgrade	103
	Unified ISSU Platform Considerations	104
	Hardware and Software Requirements Before Beginning a Unified ISSU	104
	Hardware Requirements for Unified ISSU	105
	Software Requirements for Unified ISSU	105
	Unified ISSU Terms	106
	Unified ISSU References	107
	Unified ISSU Phases Overview	107
	Unified ISSU Initialization Phase Overview	108
	Application Data Upgrade on the Standby SRP Module	109
	SNMP Traps	110
	Unified ISSU Upgrade Phase Overview	110
	Exceptions During the Upgrade Phase	111
	Verifications of Requirements	112
	Upgrade Setup	112
	Line Module Arming	113
	Line Module Control Plane Upgrade	113
	SRP Module Switchover	113
	Line Module Forwarding Plane Upgrade	114
	Unified ISSU Service Restoration Phase Overview	115
	Application Support for Unified ISSU	115
	Unexpected AAA Authentication and Authorization Behavior During Unified ISSU	124
	Unexpected ATM Behavior During Unified ISSU	124
	ILMI Sessions Not Maintained	124
	OAM CC Effects on VCC	124
	OAM VC Integrity Verification Cessation	124
	Port Data Rate Monitoring Cessation	125
	VC and VP Statistics Monitoring Halts Unified ISSU Progress	125
	Unexpected DHCP Behavior During Unified ISSU	125
	DHCP Packet Capture Halted on Line Modules	125
	Unexpected Denial-of-Service Protection Behavior During Unified ISSU	125
	Unexpected Ethernet Behavior During Unified ISSU	126
	ARP Packets Briefly Not Sent or Received	126
	Link Aggregation Interruption	126
	Port Data Rate Monitoring Halted	126
	VLAN Statistics Monitoring Halts Unified ISSU Progress	126
	Unexpected File Transfer Protocol Server Behavior During Unified ISSU	127

	IS-IS Effects on Graceful Restart and Network Stability During Unified ISSU . . .	129
	Configuring Graceful Restart Before Unified ISSU Begins	129
	Configuring Graceful Restart When BGP and LDP Are Configured	130
	Routing Around the Restarting Router to Minimize Network Instability . . .	130
	Unexpected L2TP Failover of Established Tunnels During Unified ISSU	131
	OSPF Effects on Graceful Restart and Network Stability During Unified ISSU . .	132
	Configuring Graceful Restart Before Unified ISSU Begins	132
	Configuring Graceful Restart When BGP and LDP Are Configured	132
	Configuring a Longer Dead Interval Than Normal	132
	Routing Around the Restarting Router to Minimize Network Instability . . .	133
	Unexpected Suspension of PIM During Unified ISSU	133
	Unexpected Suspension of Subscriber Login and Logouts During Unified ISSU	134
	Subscriber Statistics Accumulation or Deletion	134
	Unexpected SONET and SDH Behavior During Unified ISSU	134
	Unexpected T3 Behavior During Unified ISSU	135
	Unavailability of TACACS+ Services During Unified ISSU	135
	Interruption in Traffic Forwarding for Layer 3 Routing Protocols During Unified ISSU	135
	Recommended Settings for Routing Protocol Timers During Unified ISSU	138
	Upgrading Router Software with Unified ISSU	139
	Halt of Unified ISSU During Initialization Phase Overview	142
	Halting Unified ISSU During Initialization Phase	142
	Halt of Unified ISSU During Upgrade Phase Overview	143
	Halting Unified ISSU During Upgrade Phase	143
	Monitoring the Status of the Router During Unified ISSU	144
Chapter 6	Configuring VRRP	149
	VRRP Overview	149
	VRRP Platform Considerations	150
	VRRP Terms	151
	VRRP References	151
	VRRP Implementation in E Series Routers	152
	VRRP Router Election Rules	152
	Example: Basic VRRP Configuration	153
	Example: Commonly Used VRRP Configuration	154
	Example: VRRP Configuration Without the Real Address Owner	155
	Before You Configure VRRP	156
	Configuring VRRP	157
	Changing the Object Priority	159
	Monitoring the Configuration of VRIDs	159
	Monitoring the Configuration of VRRP Neighbors	162
	Monitoring the Statistics of VRRP Routers	163
	Monitoring the Configuration of VRRP Tracked Objects	166
Chapter 7	Managing Interchassis Redundancy	169
	ICR Overview	169
	ICR Platform Considerations	171
	Interface Specifiers	171
	ICR Terms	172

ICR References	172
ICR Scaling Considerations	172
1:1 Subscriber Redundancy in a 4–Node ICR Cluster	173
Interaction with RADIUS for ICR	173
ICR Partition Accounting Overview	175
Configuring an ICR Partition	175
Configuring the Interface on Which the ICR Partition Resides	176
Configuring VRRP Instances to Match ICR Requirements	177
Naming ICR Partitions	178
Grouping ICR Subscribers Based on S-VLAN IDs	178
Grouping ICR Subscribers Based on VLAN IDs	179
Example: Configuring ICR Partitions That Group Subscribers by S-VLAN ID . . .	180
Using RADIUS to Manage Subscribers Logging In to ICR Partitions	182
Monitoring the Configuration of an ICR Partition Attached to an Interface . . .	183
Monitoring the Configuration of ICR Partitions	184

Part 2

Index

Index	189
-------------	-----

List of Figures

Part 1	Chapters	
Chapter 1	Service Availability	3
	Figure 1: JunosE Software Service Availability Layers	4
Chapter 2	Managing Module Redundancy	9
	Figure 2: SRP Module on ERX7xx Models and ERX14xx Models	18
	Figure 3: SRP Module on the E120 and E320 Routers	19
Chapter 3	Managing Stateful SRP Switchover	35
	Figure 4: High Availability States	39
Chapter 4	Managing Stateful Line Module Switchover	67
	Figure 5: Stateful Line Module Switchover States	85
Chapter 6	Configuring VRRP	149
	Figure 6: Basic VRRP Configuration	154
	Figure 7: Commonly Used VRRP Configuration	155
	Figure 8: VRRP Configuration Without the Real Address Owner	156
Chapter 7	Managing Interchassis Redundancy	169
	Figure 9: ICR Deployment	170
	Figure 10: Sample 1:1 Subscriber Redundancy in a 4–Node ICR Cluster	173

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	9
	Table 3: Commands That Can Cause Automatic Switchover	12
	Table 4: Function of the Online and Redundant LEDs	22
	Table 5: show environment Output Fields	25
	Table 6: show hardware Output Fields	29
	Table 7: show redundancy Output Fields	31
Chapter 3	Managing Stateful SRP Switchover	35
	Table 8: Application Support for Stateful SRP Switchover	42
	Table 9: show redundancy Output Fields	57
	Table 10: show redundancy clients Output Fields	60
	Table 11: show redundancy history Output Fields	61
	Table 12: show redundancy line-card Output Fields	63
	Table 13: show redundancy srp Output Fields	64
	Table 14: show redundancy switchover-history Output Fields	65
Chapter 4	Managing Stateful Line Module Switchover	67
	Table 15: Module Configurations Supported for Stateful Switchover of LNS Sessions	70
	Table 16: show redundancy line-card slot slotNum Output Fields	96
	Table 17: show redundancy history line-card slot slotNum Output Fields	98
Chapter 5	Configuring a Unified In-Service Software Upgrade	101
	Table 18: Unified ISSU-Related Terms	106
	Table 19: Router Response to Undesirable Events During the Upgrade Phase	111
	Table 20: Application Support for Unified In-Service Software Upgrades	116
	Table 21: Behavior of Routing Protocols During a Unified In-Service Software Upgrade	136
	Table 22: Recommended Routing Protocol Timer Settings	138
	Table 23: show issu Output Fields	146
Chapter 6	Configuring VRRP	149
	Table 24: VRRP Definitions	151
	Table 25: show ip vrrp and show ip vrrp summary Output Fields	160
	Table 26: show ip vrrp neighbor Output Fields	163
	Table 27: show ip vrrp statistics Output Fields	165

Chapter 7	Table 28: show ip vrrp tracked-objects Output Fields	166
	Managing Interchassis Redundancy	169
	Table 29: ICR Terminology	172
	Table 30: show icr-partition Output Fields	183
	Table 31: show icr-partitions Output Fields	185

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 *JunosE 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 JunosE 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.	host1#show ip ospf 2 Routing Process OSPF 2 with Router ID 5.5.0.250 Router is an Area Border Router (ABR)
<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>

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

Convention	Description	Examples
(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
[] (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 } { clusterId ipAddress }

Obtaining Documentation

To obtain the most current version of all Juniper Networks technical documentation, see the Technical 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 Portable Libraries page at

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

Copies of the Management Information Bases (MIBs) for a particular software release are available for download in the software image bundle from the Juniper Networks Web site 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/us/en/local/pdf/resource-guides/7100059-en.pdf> .
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/> .
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Chapters

- Service Availability on page 3
- Managing Module Redundancy on page 9
- Managing Stateful SRP Switchover on page 35
- Managing Stateful Line Module Switchover on page 67
- Configuring a Unified In-Service Software Upgrade on page 101
- Configuring VRRP on page 149
- Managing Interchassis Redundancy on page 169

CHAPTER 1

Service Availability

This chapter explains what service availability is 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 failure, switch route processor module failure, software defects, feature upgrades, or complete router failure. These outages result in subscriber downtime.

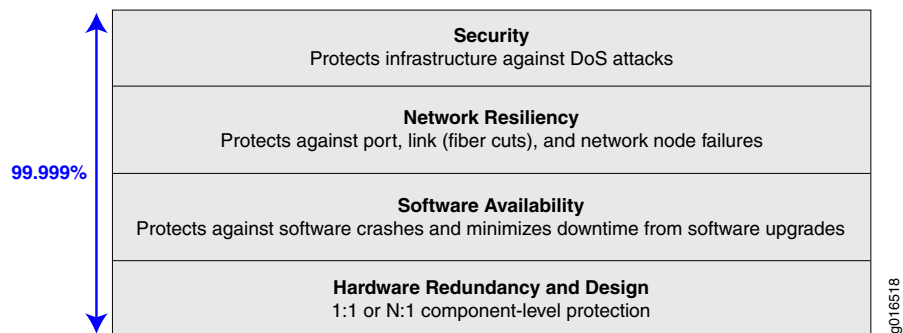
To reduce subscriber downtime, a network must have the following capabilities:

- *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 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.

JunosE 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 JunosE Software service availability.

Figure 1: JunosE Software Service Availability Layers



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 5 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 a component fails to the time when the backup component is 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 Documentation

- 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
- Stateful Line Module Switchover on page 5
- Unified ISSU on page 6
- 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.

Stateful Line Module Switchover

High availability of line modules increases the overall availability of the router by ensuring that all the subscribers who were connected during a line module recovery continue to remain logged in and can access network resources during the switchover from the primary line module to the secondary line module. Forwarding of data through the fabric slice for those subscribers continues with a brief disruption of two minutes. If you configured stateful line module switchover on a router, when a switchover occurs, a message is displayed on the active SRP module after the secondary line module successfully takes over the role of the previously configured primary line module. If the primary line module fails, the secondary line module takes the role of the primary line module. Mirrored configuration data and any mirrored volatile data are already resident in memory. The protocols and other applications re-initialize from the mirrored data and resynchronize communications with the line modules (non-volatile configuration and volatile state). Data forwarding operation continues to function normally with the secondary line module operating on behalf of the primary line module (with a small loss

of packets when the fabric is switched from the formerly active line module to the newly active line module). 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

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 router or access interface on the edge router fails. ICR accomplishes this by re-creating subscriber sessions on the backup router 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. ICR uses Virtual Router Redundancy Protocol (VRRP) to detect failures. ICR also enables you to track the failure of uplink interfaces. ICR currently supports only PPPoE subscribers.

**Related
Documentation**

- [Managing Module Redundancy on page 9](#)
- [Managing Stateful SRP Switchover on page 35](#)
- [Configuring a Unified In-Service Software Upgrade on page 101](#)
- [Configuring VRRP on page 149](#)
- [Managing Interchassis Redundancy on page 169](#)
- [Service Availability Overview on page 3](#)

CHAPTER 2

Managing Module Redundancy

This chapter describes how to manage redundancy (stateless switchover) 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 9
- Line Module Redundancy Requirements on page 10
- Understanding Automatic Switchover on page 12
- Understanding Reversion After Switchover on page 12
- Configuring Line Module Redundancy on page 13
- Managing Line Module Redundancy on page 13
- Example: Forcing the Router to Switch from Primary Line Module to Spare Line Module on page 14
- Interoperation of Redundancy and Stateful Switchover for Line Modules on page 15
- Understanding SRP Module Redundancy on page 16
- Understanding Configuration of SRP Modules for Redundancy on page 19
- Installing a Redundant SRP Module on page 20
- Managing SRP Module Redundancy on page 21
- Switching to the Redundant SRP Module on page 22
- Determination of Redundancy Status for Line Modules and SRP Modules Using Status LEDs on page 22
- Monitoring Redundancy in Installed Hardware on page 23
- Monitoring Redundancy in Line Module and SRP Modules on page 27
- Monitoring Redundancy Status on E320 Router on page 30

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*. Line modules can operate in one of the two redundancy modes: stateless switchover or high availability. Stateless switchover is the default redundancy mode. During stateless 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.



NOTE: This section does not cover behavior of line modules in high availability redundancy mode. For information about stateful line module switchover, see “Managing Stateful Line Module Switchover” on page 67.

Related Documentation

- Line Module Redundancy Requirements on page 10
- Configuring Line Module Redundancy on page 13
- Managing Line Module Redundancy on page 13
- Stateful Line Module Switchover Overview on page 68
- Guidelines for Configuring Stateful Line Module Switchover on page 70

Line Module Redundancy 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 stateless switchover redundancy for the ES2-S1 Service IOA.

IOA Behavior When the Router Reboots

On E120 and E320 routers, stateless 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 a slot 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.

Related Documentation

- Line Module Redundancy Overview on page 9
- Configuring Line Module Redundancy on page 13
- Managing Line Module Redundancy on page 13
- Stateful Line Module Switchover Overview on page 68
- Stateful Line Module Switchover Platform Considerations on page 70

Understanding Automatic Switchover

Provided you have not issued the **redundancy lockdown** 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 12).

You can also disable automatic switchover on individual slots. For more information, see “Configuring Line Module Redundancy” on page 13.

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.

Related Documentation

- Line Module Redundancy Overview on page 9
- Understanding Reversion After Switchover on page 12
- Configuring Line Module Redundancy on page 13
- Managing Line Module Redundancy on page 13

Understanding 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 13.) Before reversion can take place, the primary line module must complete the POST diagnostics.

- Related Documentation**
- Understanding Automatic Switchover on page 12
 - Line Module Redundancy Overview on page 9

Configuring Line Module Redundancy

By default, when the primary line module fails, the router automatically switches to the spare line module. Because the router is using the spare line module, no redundancy is available. The router must revert to the primary line module as soon as possible. The router does not automatically revert to the primary line module. To modify the default redundancy operations on the router, perform the following tasks:

- Disable automatic switchover on a slot.
`host1(config)#redundancy lockout 5`
- Enable automatic reversion after switchover.
`host1(config)#redundancy revertive 23:00:00`

- Related Documentation**
- Line Module Redundancy Overview on page 9
 - Line Module Redundancy Requirements on page 10
 - Managing Line Module Redundancy on page 13
 - redundancy lockout
 - redundancy revertive

Managing Line Module Redundancy

When the router is running and a redundancy group is installed, to manage stateless switchover redundancy, you must perform the following tasks:

- Force switchover manually.
`host1#redundancy force-switchover 5`
- Force reversion manually.
`host1#redundancy revert 4 23:00:00 5 September 200X`

- Related Documentation**
- Line Module Redundancy Overview on page 9
 - Line Module Redundancy Requirements on page 10
 - Configuring Line Module Redundancy on page 13
 - Stateful Line Module Switchover Overview on page 68
 - Stateful Line Module Switchover Modes on page 83

- System Operations When Stateful Line Module Switchover Is Enabled on page 74
- redundancy force-switchover
- redundancy revert

Example: Forcing the Router to Switch from Primary Line Module to Spare Line Module

In the following example, the user forces the router to switch from the primary line module to the spare line module by issuing the **redundancy force-switchover** command. In the example, you first issue the **show redundancy** command to display redundancy information specific to line modules. You can then issue the **redundancy force-switchover** command to force the router to switch from the primary line module to the spare line module. To view the status of the line modules after the switchover, you can again issue the **show redundancy** command.

1. Display the redundancy information.

```
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

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

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

2. Force the router to switch from the primary line module to the spare line module.

```
host1#redundancy force-switchover 2
```

3. Verify the redundancy status of the line module.

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

```
automatic reverting is off
```

slot	hardware role	lockout config	backed up by slot	sparing for slot	revert at
0	spare	---	---	2	---
2	primary	protected	0	---	---
12	---	---	---	---	---

slots	midplane type	midplane rev
0 - 5	6	0

Interoperation of Redundancy and Stateful Switchover for Line Modules

Line module redundancy and stateful line module switchover cannot coexist and are mutually exclusive processes. Only one of the two operations—line module redundancy or stateful line module switchover—can be enabled on a router at a point of time. You cannot configure line modules installed in a redundancy group to operate in HA mode. If a line module is a member of a redundancy group, you cannot configure that line module for stateful switchover using the **mode high-availability slot** command in Redundancy Configuration mode. Similarly, if a line module is configured for high availability, it cannot be installed in a redundancy group. If you attempt to add a line module configured in a redundancy group to the stateful switchover pair using the **mode high-availability slot** command, the particular line module is removed from the redundancy group and is added to the high availability pair. High availability configuration for a line module takes precedence over its redundancy group setting.

Related Documentation

- Stateful Line Module Switchover Overview on page 68
- mode
- line-card switch
- show redundancy line-card
- show redundancy history

Understanding SRP Module Redundancy



NOTE: This section 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 *JunosE System Basics Configuration Guide*. For information about managing high availability in a router, see “Managing Stateful SRP Switchover” on page 35.

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.



NOTE: The ERX310 router 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.

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-numbered 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 takes control without rebooting itself. For information about preventing the redundant SRP module from taking control, see “Managing SRP Module Redundancy” on page 21.

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 gives back 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 takes control, the following sequence of events occurs:

1. The original primary SRP module reboots and takes the redundant role.
2. The redundant SRP module restarts and takes the primary role without reloading new code. (When upgrading software, you must reload the software on the redundant SRP module. See *Installing JunosE Software* in the *JunosE 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 2 on page 18 and Figure 3 on page 19.)
- Issuing the **srp switch** command
- Issuing the **redundancy force-switchover** command

Figure 2: SRP Module on ERX7xx Models and ERX14xx Models





Figure 3: SRP Module on the E120 and E320 Routers



Related Documentation

- Understanding Configuration of SRP Modules for Redundancy on page 19
- Installing a Redundant SRP Module on page 20
- Managing SRP Module Redundancy on page 21

Understanding Configuration of SRP Modules for Redundancy

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 JunosE 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 *JunosE System Basics Configuration Guide*.

- Related Documentation**
- Understanding SRP Module Redundancy on page 16
 - Installing a Redundant SRP Module on page 20

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:

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.

host1#reload slot 7

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.

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



NOTE: The SRP module reboots after synchronization is complete.

Related Documentation

- Understanding SRP Module Redundancy on page 16
- Managing SRP Module Redundancy on page 21
- Switching to the Redundant SRP Module on page 22
- reload slot
- 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.

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

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

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

Related Documentation

- Understanding SRP Module Redundancy on page 16
- Installing a Redundant SRP Module on page 20
- disable-switch-on-error
- synchronize

Switching to the Redundant SRP Module

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

Switch from the primary SRP module to the redundant SRP module by doing either of the following:

- Issue the **redundancy force-switchover** command.

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

- Issue the **srp switch force** command.

```
host1#srp switch force
```

Related Documentation

- Understanding SRP Module Redundancy on page 16
- Installing a Redundant SRP Module on page 20
- Managing SRP Module Redundancy on page 21
- redundancy force-switchover
- srp switch

Determination of Redundancy Status for Line Modules and SRP Modules Using Status LEDs

You can determine the redundancy state of line modules and SRP modules by examining their status LEDs. See Table 4 on page 22 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.

Related Documentation

- Line Module Redundancy Overview on page 9
- Understanding SRP Module Redundancy on page 16

Monitoring Redundancy in Installed Hardware

Purpose Display redundancy information specific to the hardware installed.

Action To display redundancy information of the hardware installed, system environment information, and detailed information on the temperature status of an ERX7xx router.

```
host1#show environment all
  chassis: 14 slot (id 0x3, rev. 0x0)
  fabric: 5 Gbps (rev. 1)
  fans: ok
  nvs: ok (81MB flash disk, 54% full)
  power: A ok, B not present
  AC power: A not present, B not present
  srp redundancy: none
*** slots: cards missing or offline
      online: 6 9
      standby: 8
      offline: 2
      empty: 0 1 3 4 5 7 10 11 12 13
  line redundancy: 1 redundancy group(s)
      width 6, spare 8, primary 9
  temperature: ok
  timing: primary
      primary: internal SC oscillator (ok)
      secondary: internal SC oscillator (ok)
      tertiary: internal SC oscillator (ok)
      auto-upgrade enabled
*** system operational: no
```

slot	processor temperature (10C - 70C)	processor temperature status	IOA temperature (10C - 70C)	IOA temperature status
0	31	normal	30	normal
3	31	normal	30	normal
5	31	normal	30	normal
7	31	normal	30	normal

```
processor temperature ranges
  below -5C is too cold
  above 80C is too hot
  low temperature warning below 10C
  high temperature warning above 70C
IOA temperature ranges
  below -5C is too cold
  above 80C is too hot
  low temperature warning below 10C
  high temperature warning above 70C
```

To display redundancy information of the hardware installed, system environment information, and detailed information on the temperature status of an E320 router.

```
host1#show environment all
  chassis: 17 slot (id 0x3, rev. 0x0)
  fabric: 100 Gbps (rev. 1)
  fans: fanSubsystemOk
  nvs: ok (977MB flash disk, 29% full), matches running config
  power: A ok, B not present
  srp redundancy: mode is file-system-synchronization      auto-sync
```

```

enabled, switch-on-error enabled
status unknown
*** slots: cards missing or offline
online: 0 6 13
offline: 7
empty: 1 2 3 4 5 11 12 14 15 16
fabric slots: ok
online: 6 7 8 9 10
line redundancy: none
temperature: ok
timing: primary
primary: internal SC oscillator (ok)
secondary: internal SC oscillator (ok)
tertiary: internal SC oscillator (ok)
auto-upgrade enabled
fabric redundancy: ok

*** system operational: no

```

slot	type	temperature (10C - 70C)	temperature status
0	LM-4	42	normal
0/1	GE-4 IOA	23	normal
6	SRP-100	32	normal
6	SFM-100	32	normal
6/0	SRP IOA	25	normal
7	SFM-100	30	normal
8	SFM-100	23	normal
9	SFM-100	25	normal
10	SFM-100	24	normal
13	LM-4	24	normal
13/0	GE-4 IOA	23	normal

```

fabric temperature ranges
  below -5C is too cold
  above 79C is too hot
  low temperature warning below 10C
  high temperature warning above 70C
processor temperature ranges
  below -5C is too cold
  above 79C is too hot
  low temperature warning below 10C
  high temperature warning above 70C
IOA temperature ranges
  below -5C is too cold
  above 79C is too hot
  low temperature warning below 10C
  high temperature warning above 70C

```

Meaning Table 5 on page 25 lists the **show environment** command output fields.

Table 5: show environment Output Fields

Field Name	Field Description
Chassis	<p>Number of slots, midplane identifier, and hardware revision number:</p> <ul style="list-style-type: none"> • 14Slot—5Gbps, 14 slot midplane • midplaneId7Slot—5Gbps, 7 slot midplane • midplaneIdRx1400—10Gbps ASIC compatible, 12 line module slots, 2 SRP module slots for ERX14xx models • midplaneIdRx700—10Gbps ASIC compatible, 5 line module slots, 2 SRP module slots for ERX7xx models • 17Slot—100 or 320 Gbps, 17-slot midplane for the E120 router • 11Slot—320Gbps, 11-slot midplane for the E120 router
fabric	Capacity and hardware revision of the fabric.
fans	Status of fans.
nvs	Status and capacity of NVS and amount of space used.
power	Status of power feeds.
AC power	For ERX310 router only; status of power feeds.
srp redundancy	Availability of a redundant SRP module
slots: cards missing or offline.	<p>Status of each slot:</p> <ul style="list-style-type: none"> • online • standby • offline • empty
line redundancy	<p>Number of redundancy groups installed:</p> <ul style="list-style-type: none"> • width—Number of slots the redundant midplane covers • spare—Slot that contains a spare line module • primary—Slot that contains the primary line module
fabric redundancy	Status of redundancy on the switch fabric on the E120 and E320 routers. Possible values: ok and none.
temperature	Status of the system temperature

Table 5: show environment Output Fields (*continued*)

Field Name	Field Description
timing	Source of the timing signal: <ul style="list-style-type: none"> primary—Type and status of the primary timing signal secondary—Type and status of the secondary timing signal tertiary—Type and status of the tertiary timing signal auto-upgrade—Status of the auto-upgrade parameter, which enables the system to revert to a higher-priority timing source after switching to a lower-priority timing source.
system operational	Status of the system
slot	Number of the slot in which the module resides
type	Type of module in the slot on the E120 and E320 routers
temperature	Temperature of the line module, SRP module, or SFM on the E120 and E320 routers
processor temperature	Temperature of the line module or SRP module
processor temperature status	Temperature condition of the line module: <ul style="list-style-type: none"> normal—Temperature is within the normal range too hot—Module is too hot; system will go into thermal protection mode if the temperature of any module exceeds 80 C too cold—Module is too cold; system will go into thermal protection mode if the temperature of any module drops below -5 C
IOA temperature	Temperature of the corresponding I/O module or IOA
IOA temperature status	Temperature condition of the corresponding I/O module or IOA: <ul style="list-style-type: none"> normal—Temperature is within the normal range too hot—Module is too hot; system will go into thermal protection mode if the temperature of any module exceeds 80 C too cold—Module is too cold; system will go into thermal protection mode if the temperature of any module drops below -5 C
processor temperature ranges	Temperature ranges for the line modules and SRP modules.

Table 5: show environment Output Fields (*continued*)

Field Name	Field Description
IOA temperature ranges	Temperature ranges for the I/O modules on ERX7xx models, ERX14xx models, and the ERX310 router or IOAs on the E120 and E320 routers.
fabric temperature ranges	Temperature ranges for the SRP modules and SFMs on the E120 and E320 routers.

Related Documentation • show environment

Monitoring Redundancy in Line Module and SRP Modules

Purpose Display redundancy information about SRP modules, line modules, and I/O modules in ERX7xx models, ERX14xx models, and the ERX310 router. Also displays redundancy information about SRP modules, line modules, and IOAs in the E120 router and the E320 router.

Action To display redundancy information about the SRP modules, line modules, and I/O modules on an ERX7xx router.

host1#show hardware

slot	type	serial number	assembly number	assembly rev.	ram (MB)
0	SRP-10Ge	4305358981	3500005472	A06	2048
1	SRP-10Ge	4305359020	3500005472	A06	2048
2	---	---	---	---	---
3	---	---	---	---	---
4	CT3-12	4305337201	3500010901	A07	128
5	OC3/OC12/DS3-ATM	4605300290	3500103958	A06	256
6	GE/FE	4605340294	3500104554	A08	256

slot	type	serial number	assembly number	assembly rev.	number of MAC addresses
0	---	---	---	---	---
1	SRP-10Ge I/O	4605250426	3500003302	A02	1
2	---	---	---	---	---
3	---	---	---	---	---
4	CT3/T3-12 I/O	4305316605	3500010801	A02	---
5	OC3(8)-MM I/O	4304443600	4500001501	A03	4
6	GE-SFP I/O	4605310064	4500002001	A05	1

slot	base MAC address
0	---
1	0090.1aa0.577a
2	---
3	---
4	---

```

5      0090.1a41.7c68
6      0090.1aa0.6216

```

To display redundancy information about the SRP modules, line modules, and IOAs on the E320 router.

```
host1#show hardware
```

```

Chassis
-----
type      serial      assembly      assembly      Major/Minor
number    number      rev.          rev
-----
Chassis   5504200687  4400006402    01            0.101

Modules
-----
slot      type      serial      assembly      assembly      ram      Major/Min
number    number    number      rev.          (MB)         rev
-----
0          ---      ---         ---           ---          ---      ---
1          ---      ---         ---           ---          ---      ---
2          LM-4      4303470363  4500006301    01           256      1.101
3          ---      ---         ---           ---          ---      ---
4          ---      ---         ---           ---          ---      ---
5          ---      ---         ---           ---          ---      ---
6          ---      ---         ---           ---          ---      ---
6          ---      ---         ---           ---          ---      ---
7          SRP-100   4304218323  4500006601    03           1024     1.103
7          SFM-100   4304218323  4500006601    03           ---      1.103
8          SFM-100   4304206756  4500006701    04           ---      1.104
9          SFM-100   4304206762  4500006701    04           ---      1.104
10         SFM-100   4304206737  4500006701    04           ---      1.104
11         ---      ---         ---           ---          ---      ---
12         ---      ---         ---           ---          ---      ---
13         ---      ---         ---           ---          ---      ---
14         ---      ---         ---           ---          ---      ---
15         ---      ---         ---           ---          ---      ---
16         ---      ---         ---           ---          ---      ---

Adapters
-----
slot      type      serial      assembly      assembly      number
number    type      number      number      rev.          of
          type      number      number      rev.          MAC
          type      number      number      rev.          addresses
-----
0/0        ---      ---         ---           ---          ---
0/1        ---      ---         ---           ---          ---
1/0        ---      ---         ---           ---          ---
1/1        ---      ---         ---           ---          ---
2/0        GE-4 IOA  4304020462  4500006800    11           4
2/1        ---      ---         ---           ---          ---
3/0        ---      ---         ---           ---          ---
3/1        ---      ---         ---           ---          ---
4/0        ---      ---         ---           ---          ---
4/1        ---      ---         ---           ---          ---
5/0        ---      ---         ---           ---          ---
5/1        ---      ---         ---           ---          ---
7/0        SRP IOA   4303470366  4500006500    02           2
11/0       ---      ---         ---           ---          ---
11/1       ---      ---         ---           ---          ---

```

12/0	---	---	---	---	---
12/1	---	---	---	---	---
13/0	---	---	---	---	---
13/1	---	---	---	---	---
14/0	---	---	---	---	---
14/1	---	---	---	---	---
15/0	---	---	---	---	---
15/1	---	---	---	---	---
16/0	---	---	---	---	---
16/1	---	---	---	---	---
slot	base MAC address	Major/Minor rev			
0/0	---	---			
0/1	---	---			
1/0	---	---			
1/1	---	---			
2/0	0090.1a00.17ec	1.111			
2/1	---	---			
3/0	---	---			
3/1	---	---			
4/0	---	---			
4/1	---	---			
5/0	---	---			
5/1	---	---			
7/0	0090.1a00.17ae	1.102			
11/0	---	---			
11/1	---	---			
12/0	---	---			
12/1	---	---			
13/0	---	---			
13/1	---	---			
14/0	---	---			
14/1	---	---			
15/0	---	---			
15/1	---	---			
16/0	---	---			
16/1	---	---			
Fan(s)					

Tray	type	serial number	assembly number	assembly rev.	Major/Minor rev
0	Primary FAN	4303370009	4400007000	01	1.101

Meaning Table 6 on page 29 lists the **show hardware** command output fields.

Table 6: show hardware Output Fields

Field Name	Field Description
Slot	Physical slot that contains the module.
type	Kind of module or chassis and fan tray in the E120 and E320 routers; an "e" at the end of an SRP module type (for example, SRP-5Ge) indicates that the module includes error-checking code (ECC) memory.

Field Name	Field Description
serial number	Serial number of the module, chassis, or fan tray.
assembly number	Part number of the module, chassis, or fan tray.
assembly rev.	Hardware revision of the module, chassis, or fan tray.
ram (MB)	Memory capacity of the host processor.
number of MAC addresses	Total number of Ethernet addresses on an I/O module or an IOA.
base MAC address	Lowest Ethernet address on an I/O module or an IOA
Tray	Number of the fan tray in the E120 and E320 routers; 0 indicates the primary fan.
Major/Minor rev	Revision number of the module on the E120 and E320 routers.

Monitoring Redundancy Status on E320 Router

Action host1#show redundancy

Criteria Preventing High Availability from being Active

```
automatic reverting is off
```

Copyright © 2010, Juniper Networks, Inc.

```

-----
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     ---      ---

```

Meaning Table 7 on page 31 lists the **show redundancy** command output fields.

Table 7: show redundancy Output Fields

Field Name	Field Description
SRP	
high-availability state	<p>State of high availability mode:</p> <ul style="list-style-type: none"> disabled—Initial, default state for high-availability mode. The router continues to use file system synchronization. active—Data synchronized from the active SRP module to the standby SRP module during initialization remains synchronized through mirroring updates. pending—If an unsupported application is configured, the router transitions to this state. initializing—If SRP module is in initializing state, bulk synchronization of memory and NVS occurs.
current redundancy mode	<p>Redundancy mode currently used by the router:</p> <ul style="list-style-type: none"> high-availability—Ensures rapid SRP module recovery after a switchover by using initial bulk file transfer and subsequent, transaction-based mirroring. file-system-synchronization—Default redundancy mode of the router. SRP modules reload all line modules and restart from saved configuration files.
last activation type	<p>Last type of activation that occurred on the router. The method using which the SRP last booted:</p> <ul style="list-style-type: none"> cold-start—When the router is in pending state and switchover occurs, the router undergoes a cold-start or cold switch. warm-start—When the router is in active state and switchover occurs, the router undergoes a warm switch or warm-start.

Table 7: show redundancy Output Fields (*continued*)

Field Name	Field Description
Criteria Preventing High Availability from being Active	Criteria preventing the router from being in the active state of high availability mode. NOTE: For the router to be in the Active state, all criteria for this option must be “yes”.
Criteria Required for High Availability to be Active	Criteria required for the router to be in the active state of high availability mode. NOTE: For the router to be in the Active state, all criteria for this option must be “yes”.
Line Card	
automatic reverting	State of automatic reverting. Possible states: on or off.
slots	Slots in which the line modules reside.
hardware role	Function of the line module. Possible values: primary or spare.
lockout config	Status of redundancy on the line module: <ul style="list-style-type: none"> 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 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.
slice state	State of the fabric slice. Possible values: online or not present.
type	Identifier for the type of hardware. Possible values: SRP modules or SFM modules.

- Related Documentation**
- [show redundancy](#)
 - [show redundancy line-card](#)
 - [show redundancy srp](#)

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 E Series routers. Use this chapter with “Managing Module Redundancy” on page 9 to fully manage the SRP features.

This chapter contains the following sections:

- Stateful SRP Switchover Overview on page 35
- Stateful SRP Switchover Platform Considerations on page 36
- Stateful SRP Switchover Redundancy Modes on page 37
- Stateful SRP Switchover States on page 38
- Application Support for Stateful SRP Switchover on page 42
- Guidelines for Activating High Availability on page 51
- Activating High Availability on page 52
- Guidelines for Deactivating High Availability on page 53
- Deactivating High Availability on page 53
- Guidelines for Setting the IP Interface Priority on page 54
- Setting the IP Interface Priority on page 54
- Guidelines for Upgrading Software on page 55
- Monitoring the Redundancy Status on page 56
- Monitoring the Redundancy Status of Applications on page 59
- Monitoring the Redundancy History on page 61
- Monitoring the Redundancy Status of Line Modules on page 62
- Monitoring the Redundancy Status of SRP Modules on page 63
- Monitoring the Redundancy Switchover History on page 65
- Clearing the Redundancy History on page 65

Stateful SRP Switchover Overview

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*.

Related Documentation

- Stateful SRP Switchover Redundancy Modes on page 37
- Stateful SRP Switchover States on page 38
- Application Support for Stateful SRP Switchover on page 42

Stateful SRP Switchover 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.

Related Documentation

- Monitoring the Redundancy Status on page 56
- Monitoring the Redundancy Status of SRP Modules on page 63
- `show redundancy`
- `show redundancy srp`

Stateful SRP Switchover Redundancy Modes

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 JunosE 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 9.

High Availability Mode

Currently applicable to the SRP module, 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.

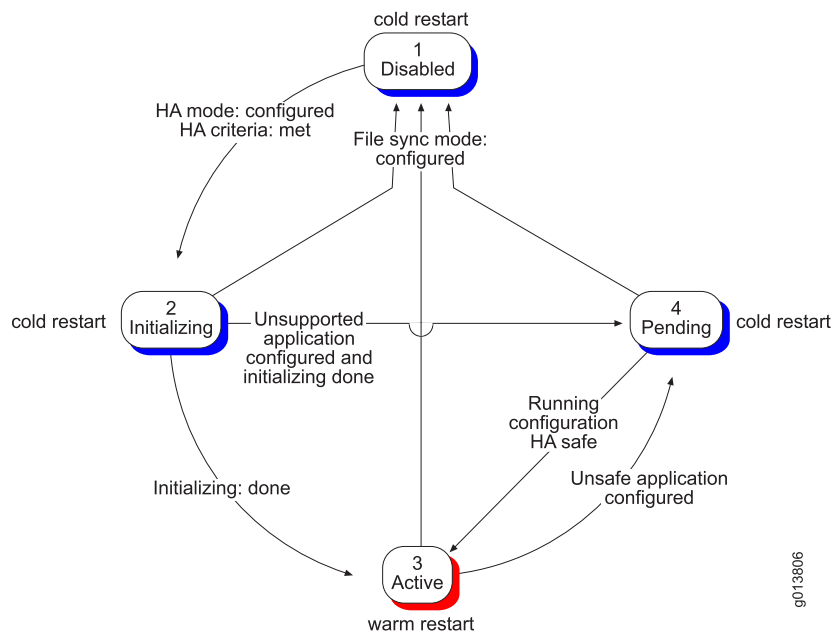
Related Documentation

- Stateful SRP Switchover States on page 38
- `disable-autosync`
- `redundancy`
- `mode`

Stateful SRP Switchover States

The SRP progresses through various high availability states. These states are illustrated in Figure 4 on page 39.

Figure 4: High Availability States



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.

Related Documentation

- Monitoring the Redundancy Status on page 56
- Monitoring the Redundancy Status of SRP Modules on page 63

- show redundancy
- show redundancy srp
- synchronize

Application Support for Stateful SRP Switchover

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 8 on page 42 indicates which applications support or do not support stateful SRP switchover.

Application Support

Table 8: Application Support for Stateful SRP Switchover

Application	Supported	Unsupported	Notes
Physical Layer Protocols			
DS1	✓	—	—
DS3	✓	—	—
HDLC	✓	—	—
SONET/SDH	✓	—	—
SONET/SDH VT	✓	—	—
Link-Layer Protocols			

Table 8: Application Support for Stateful SRP Switchover (*continued*)

Application	Supported	Unsupported	Notes
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	✓	–	–
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.

Table 8: Application Support for Stateful SRP Switchover (*continued*)

Application	Supported	Unsupported	Notes
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.
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			

Table 8: Application Support for Stateful SRP Switchover (*continued*)

Application	Supported	Unsupported	Notes
Multicast Routing	✓	—	Stateful SRP switchover. During switchover, the system mirrors the multicast queue so that IP can use the same queue without needing to recreate a different connection. The multicast queues are also preserved during the switchover and graceful restart period to ensure that multicast data continues to be forwarded using the previously learned multicast forwarding state.
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 MGMTM that graceful restart is complete.</p>

Table 8: 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>Stateful SRP switchover. During switchover, the system mirrors the multicast queue so that IP can use the same queue without needing to recreate a different connection. The multicast queues are also preserved during the switchover and graceful restart period to ensure that multicast data continues to be forwarded using the previously learned multicast forwarding state.</p>

Table 8: 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 8: 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.

Table 8: Application Support for Stateful SRP Switchover (*continued*)

Application	Supported	Unsupported	Notes
Remote Access			
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 now supports stateful SRP switchover (high availability).</p> <p>After SRP warm switchover, the router restores the client bindings from the mirrored DHCPv6 information as it does for other applications that support stateful SRP switchover.</p>
L2TP	✓	–	–
L2TP Dialout	–	✓	–

Table 8: Application Support for Stateful SRP Switchover (*continued*)

Application	Supported	Unsupported	Notes
IPv4 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 IPv4 local address server having full HA support.
IPv6 Local Address Pools	✓	—	When the IPv6 local pools are configured, you can perform an HA switchover without cold booting the router because the configuration is now HA safe. The prefix assigned to the subscriber, before and after the warm restart, remains the same. The In Use prefix count also remains the same before and after the warm restart.
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.

Table 8: Application Support for Stateful SRP Switchover (*continued*)

Application	Supported	Unsupported	Notes
J-Flow (IP flow statistics)	✓	–	–
Line Module Redundancy	✓	–	–
Network Address Translation	✓	–	–
NTP	✓	–	–
Resource Threshold Monitor	✓	–	–
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.

Related Documentation

- Monitoring the Redundancy Status of Applications on page 59
- show redundancy clients

Guidelines for Activating High Availability

Before you activate high availability on the SRP modules, you must be aware of any high availability–related changes to SRP management commands. For information on high availability–related changes to SRP, see “Managing Stateful SRP Switchover” on page 35.

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

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.

**Related
Documentation**

- Stateful SRP Switchover Redundancy Modes on page 37
- Activating High Availability on page 52
- redundancy
- mode

Activating High Availability

The switch route processor (SRP) module can operate in one of the two redundancy modes—file system synchronization and high availability. When you activate high availability, 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.

To activate high availability:

1. From Global Configuration mode, launch Redundancy Configuration mode.
`host1(config)#redundancy`
2. In Redundancy Configuration mode, specify high availability as the redundancy mode.
`host1(config-redundancy)#mode high-availability`

**Related
Documentation**

- Stateful SRP Switchover Redundancy Modes on page 37
- Guidelines for Activating High Availability on page 51
- redundancy
- mode

Guidelines for Deactivating High Availability

You can disable high availability (stateful SRP switchover) by launching Redundancy Configuration mode and issuing the **mode file-system-synchronization** command or specifying the **no mode** command.

In the **file-system-synchronization** mode, the router synchronizes the files and data such as configuration files and releases that are stored in NVS (nonvolatile storage) between the primary and standby SRP modules. This is the default behavior mode for E Series routers that contain redundant SRPs.

Because this mode uses file synchronization instead of transaction-based mirroring, when the active SRP module switches to the standby SRP, the router cold-starts.

Related Documentation

- Stateful SRP Switchover Redundancy Modes on page 37
- Deactivating High Availability on page 53
- redundancy
- mode

Deactivating High Availability

The switch route processor (SRP) module can operate in one of the two redundancy modes—file system synchronization and high availability. When you disable high availability, the router uses file system synchronization mode which is the default behavior mode for E Series routers that use redundant SRPs. The router synchronizes the contents of the NVS (nonvolatile storage) in the primary and standby SRP modules.

To disable high availability support:

1. From Global Configuration mode, launch Redundancy Configuration mode.
`host1(config)#redundancy`
2. In Redundancy Configuration mode, you can disable high availability by doing one of the following:
 - Specify file system synchronization mode as the redundancy mode.
`host1(config-redundancy)#mode file-system-synchronization`
 - Specify the **no** version to disable high availability.
`host1(config-redundancy)#no mode`

Related Documentation

- Stateful SRP Switchover Redundancy Modes on page 37
- Guidelines for Deactivating High Availability on page 53
- redundancy
- mode

Guidelines for 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.

Related Documentation

- Setting the IP Interface Priority on page 54
- `ip initial-sequence-preference`
- `ipv6 initial-sequence-preference`

Setting the IP Interface Priority

Use the **ip initial-sequence-preference** command to set the preference value on an IP or IPv6 interface at the subinterface level. To configure the interface as high-priority, specify the value of the initial sequence preference as 1. To configure the interface as low-priority, specify the value as 0.

To set the priority for the IPv4 or IPv6 interface, you can do one of the following:

- From Subinterface Configuration mode, explicitly configure the IPv4 interface as high-priority:

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

- From Subinterface Configuration mode, explicitly configure the IPv6 interface as low-priority:

```
host1(config-subif)#ipv6 initial-sequence-preference 0
```

- Related Documentation**
- Guidelines for Setting the IP Interface Priority on page 54
 - ip initial-sequence-preference
 - ipv6 initial-sequence-preference

Guidelines for 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).

- Related Documentation**
- Stateful SRP Switchover Redundancy Modes on page 37
 - Stateful SRP Switchover States on page 38
 - srp switch

Monitoring the Redundancy Status

Purpose Display the redundancy modes and other information about stateful SRP switchover.

Action To display summary redundancy status information.

```
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

```
-----
              criterion                               met
-----
Standby SRP is online and capable of mirroring?      No
```

```
Line Card
```

```
-----
```

```
automatic reverting is off
```

```

                                backed
                                up
                                by
slot  hardware  lockout  slot  sparing  revert
-----
3      ---      ---      ---      ---      ---
8      spare      ---      ---      ---      ---
12     primary  protected ---      ---      ---
```

```

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

To display detailed redundancy status information:

```
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

```
-----
              criterion                               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

                                backed
                                up
                                by
slot  hardware  lockout  slot  sparing  revert
-----
3      ---      ---      ---      ---      ---
8      spare    ---      ---      ---      ---
12     primary  protected ---      ---      ---

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

Meaning Table 9 on page 57 lists the **show redundancy** command output fields.

Table 9: show redundancy Output Fields

Field Name	Field Description
SRP	
high-availability state	State of high availability mode: <ul style="list-style-type: none">disabled—Initial, default state for high-availability mode. The router continues to use file system synchronization.active—Data synchronized from the active SRP module to the standby SRP module during initialization remains synchronized through mirroring updates.pending—If an unsupported application is configured, the router transitions to this state.initializing—If SRP module is in initializing state, bulk synchronization of memory and NVS occurs.
current redundancy mode	Redundancy mode currently used by the router: <ul style="list-style-type: none">high-availability—Ensures rapid SRP module recovery after a switchover by using initial bulk file transfer and subsequent, transaction-based mirroring.file-system-synchronization—Default redundancy mode of the router. SRP modules reload all line modules and restart from saved configuration files.

Table 9: show redundancy Output Fields (*continued*)

Field Name	Field Description
last activation type	<p>Last type of activation that occurred on the router. The method using which the SRP last booted:</p> <ul style="list-style-type: none"> cold-switch—When the router is in pending state and switchover occurs, the router undergoes a cold-switch or cold re-start. warm-switch—When the router is in active state and switchover occurs, the router undergoes a warm-switch or warm re-start.
Criteria Preventing High Availability from being Active	<p>Criteria preventing the router from being in the active state of high availability mode.</p> <p>NOTE: For the router to be in the Active state, all criteria for this option must be “yes”.</p>
Criteria Required for High Availability to be Active	<p>Criteria required for the router to be in the active state of high availability mode.</p> <p>NOTE: For the router to be in the Active state, all criteria for this option must be “yes”.</p>
Line Card	
automatic reverting	State of automatic reverting. Possible states: on or off.
slots	Slots in which the line modules reside.
hardware role	Function of the line module. Possible values: primary or spare.
lockout config	<p>Status of redundancy on the line module:</p> <ul style="list-style-type: none"> 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 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.

Table 9: show redundancy Output Fields (*continued*)

Field Name	Field Description
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. Possible values: online or not present.
type	Identifier for the type of hardware. Possible values: SRP modules or SFM modules.

Related Documentation

- [show redundancy](#)

Monitoring the Redundancy Status of Applications

Purpose Display the redundancy status of the applications.

Action To display the applications that do not support high availability.

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

You can also display the redundancy status information of all clients. Specifies whether the client supports high availability and also the safety level of configuration. For instance, if an unsupported client is configured on a router with high availability enabled, the configuration reads "unsafe".

host1#show redundancy clients all

High Availability Client Information

client	mode	configuration
atm1483DataService	supported	safe
AA83	supported	safe
aaaServer	supported	safe
atmAa15	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
s1otHelper	supported	safe
Cisco HDLC	supported	safe
ServiceManager	supported	safe
Sonet	supported	safe
SonetPath	supported	safe
SonetVT	supported	safe
IPsec Tunnel (ST)	supported	safe

Meaning Table 10 on page 60 lists the **show redundancy clients** command output fields.

Table 10: show redundancy clients Output Fields

Field Name	Field Description
client	High availability client.
mode	High availability status of the client. Possible values: supported or unsupported.
Configuration	Safety level of the configuration based on whether or not the client is supported or unsupported and in case of those unsupported, whether or not the client has been configured. For example, if an unsupported client has been configured on a router with high availability enabled, the configuration reads "unsafe".

Related Documentation

- show redundancy clients

Monitoring the Redundancy History

Purpose Display information about dates, times, and the number of occurrences for starts and switchovers.

Action To display information about the number of occurrences for starts and switchovers.

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
```

To display the additional redundancy history information:

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

Meaning Table 11 on page 61 lists the **show redundancy history** command output fields.

Table 11: show redundancy history Output Fields

Field Name	Field Description
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.

Table 11: show redundancy history Output Fields (*continued*)

Field Name	Field Description
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.
SRP activation time	Amount of time the SRP module has been active.
type	Last type of activation that occurred on this router.
slot	Slot in which the line module resides.
system uptime	Amount of time the chassis has been operational.
running release	Release running on the SRP module at the time.

Related Documentation

- show redundancy history

Monitoring the Redundancy Status of Line Modules

Purpose Display redundancy information specific to line modules.

Action To display the redundancy status of the line modules.

```
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			

Meaning Table 12 on page 63 lists the **show redundancy line-card** command output fields.

Table 12: show redundancy line-card Output Fields

Field Name	Field Description
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: <ul style="list-style-type: none"> 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.

Related Documentation

- show redundancy line-card

Monitoring the Redundancy Status of SRP Modules

Purpose Display redundancy information specific to SRP modules.

Action To display the redundancy status of the SRP modules.

```
host1#show redundancy srp
```

```
high-availability state: active
current redundancy mode: high-availability
last activation type: warm-switch
```

To display the redundancy status of the SRP modules in detail.

```
host1#show redundancy srp detail
```

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

Criteria Required for High Availability to be Active

```
-----
              criterion                               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

Meaning Table 13 on page 64 lists the **show redundancy srp** command output fields.

Table 13: show redundancy srp Output Fields

Field Name	Field Description
high-availability state	<p>State of high availability mode:</p> <ul style="list-style-type: none"> disabled—Initial, default state for high-availability mode. The router continues to use file system synchronization. active—Data synchronized from the active SRP module to the standby SRP module during initialization remains synchronized through mirroring updates. pending—If an unsupported application is configured, the router transitions to this state. initializing—If SRP module is in initializing state, bulk synchronization of memory and NVS occurs.
current redundancy mode	<p>Redundancy mode currently being used by this router:</p> <ul style="list-style-type: none"> high-availability—Ensures rapid SRP module recovery after a switchover by using initial bulk file transfer and subsequent, transaction-based mirroring. file-system-synchronization—Default redundancy mode of the router. SRP modules reload all line modules and restart from saved configuration files.
last activation type	<p>Last type of activation that occurred on the router. The method using which the SRP last booted:</p> <ul style="list-style-type: none"> cold-switch—When the router is in pending state and switchover occurs, the router undergoes a cold-switch or cold re-start. warm-switch—When the router is in active state and switchover occurs, the router undergoes a warm-switch or warm re-start.
Criteria Required for High Availability to be Active	<p>Criteria required for high availability to be active.</p> <p>NOTE: All criteria must be “yes” for high availability to be active.</p>

Related Documentation

- show redundancy srp

Monitoring the Redundancy Switchover History

Purpose Display information about stateful SRP switchover history for the chassis.

Action `host1# show redundancy switchover-history`

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
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

Meaning Table 14 on page 65 lists the `show redundancy switchover-history` command output fields.

Table 14: show redundancy switchover-history Output Fields

Field Name	Field Description
SRP activation time	Amount of time the SRP module has been active.
type	Type of switchover.
slot	Slot in which the SRP module resides.
system uptime	Amount of time the chassis has been operational.
running release	Release running on the SRP module at the time of the switchover.

Related Documentation

- `show redundancy switchover-history`

Clearing the Redundancy History

To clear the stateful SRP switchover history for the router:

- Issue the `clear redundancy history` command:

```
host1#clear redundancy history
```

There is no **no** version.

- Related Documentation**
- [Monitoring the Redundancy History on page 61](#)
 - [Monitoring the Redundancy Switchover History on page 65](#)
 - [clear redundancy history](#)
 - [show redundancy history](#)
 - [show redundancy switchover-history](#)

CHAPTER 4

Managing Stateful Line Module Switchover

This chapter describes how to manage stateful line module switchover (high availability) software features for E120 and E320 routers, and contains the following sections:

- Stateful Line Module Switchover Overview on page 68
- Benefits of Stateful Line Module Switchover on page 68
- Stateful Line Module Switchover Platform Considerations on page 70
- Guidelines for Configuring Stateful Line Module Switchover on page 70
- System Operations When Stateful Line Module Switchover Is Enabled on page 74
- Stateful Line Module Configuration Scenarios on page 75
- Replacement of Line Modules When Stateful Line Module Switchover Is Enabled on page 77
- Application Support for Stateful Line Module Switchover on page 79
- Stateful Line Module Switchover Modes on page 83
- Stateful Line Module Switchover States on page 84
- Guidelines for Activating High Availability on page 87
- Activating High Availability on page 88
- Guidelines for Deactivating High Availability on page 89
- Deactivating High Availability on page 89
- Switching Over from a Primary Line Module to Secondary Line Module on page 90
- Log Messages Generated for Stateful LM Switchover on page 91
- Preservation of Statistics During Stateful Line Module Switchover on page 93
- Performance Impact and Scalability Considerations on page 94
- Use of Status LEDs to Monitor the High Availability States of Line Modules on page 95
- Monitoring the Redundancy Status of Line Modules in a Specific Slot on page 95
- Monitoring the Redundancy History of Line Modules in a Specific Slot on page 97

Stateful Line Module Switchover Overview

JunosE Software now supports high availability for ES2 4G line modules configured with Service IOAs on E120 and E320 routers. These line modules function in a 1:1 redundancy mode with the active module as the primary line module and the spare or standby module as the secondary line module. This functionality of high availability for line modules is also referred to as *stateful line module switchover*.

In releases in which the stateful line module switchover feature was not supported or in scenarios in which this behavior is disabled, the restart of a line module causes it to be reloaded and all the subscriber sessions that are routed through it to be disconnected. All users connected to the router when the line module is reloading need to log in again and reestablish their connections. Based on the router models deployed in your environment, the configuration settings applied in your network, and the number of active subscriber sessions, reestablishment and resetting of subscriber connections can consume several minutes.

Stateful line module switchover reduces the impact on subscriber traffic during a stateful switchover from the active line module to the standby line module by ensuring that existing subscriber sessions remain active with a brief disconnection in traffic. Stateful line module switchover maintains user sessions and reduces data traffic outage through the router to a brief duration during the switchover, thereby improving the overall availability of the router. Stateful line module switchover keeps the connections through the module up, with a brief disruption in forwarding on existing interfaces through the fabric slice. When you restart the line module, applications recover to a stable state by synchronizing with their peers on the SRP module swiftly to resume normal operations. This mechanism enables the system to keep user connections up and data forwarding outage minimal through the fabric slice.

- Related Documentation**
- Benefits of Stateful Line Module Switchover on page 68
 - Stateful Line Module Switchover Platform Considerations on page 70

Benefits of Stateful Line Module Switchover

Line module high availability enhances the overall reliability and resiliency of the router. All subscribers who are connected during line module recovery remain active; their sessions are preserved during switchover and forwarding of data through the fabric slice is disrupted briefly. When a switchover occurs with the high availability state being active, a message is displayed on the active SRP module after the secondary line module has successfully taken over from the previously configured primary line module.

If you configured the ha logging event category (using the **log severity notice ha** command), log messages are displayed when high availability for line modules is enabled or disabled, either because of manual settings or in response to system events. For example, the following log message is displayed when a line module event causes high availability to be disabled.

NOTICE 01/27/2004 19:54:06 ha (): High Availability disabled due to secondary line card in slot 9 down

The commands used to configure stateful switchover for SRP modules and line modules are very similar. The configuration modes are the same. The keywords to configure stateful switchover for SRP modules and for line modules are different. When redundancy mode is configured for high availability on E120 and E320 routers with dual SRP modules, high availability is enabled only for the corresponding SRP modules. You must explicitly specify high availability for the line modules to enable stateful switchover of the line modules to occur.

1:1 Redundancy Model

Line module high availability uses a 1:1 redundancy model to maintain subscriber sessions, during and after a switchover of the primary line module to the secondary line module. This feature is supported only on E120 and E320 routers installed with ES2 4G LMs and Service IOAs. You can enable line module high availability feature only on the compatible line modules and IOA combinations. On routers in which line module high availability is disabled or not available for configuration (cold-restart support is not present), all subscriber sessions are forcibly terminated when a line module fails or stops responding. This mechanism is known as a cold-restart.

Seamless Preservation of Subscriber Sessions

When the secondary line module in the HA configuration becomes the primary line module, all the applications on it recover to a stable state to operate normally in the place of the failed line module. The new primary line module maintains existing subscriber sessions and continues to forward subscriber data traffic after the switchover. A brief subscriber data disconnection occurs for two minutes during the switchover. Diagnostic functions are still run on the failed line module to ensure that the hardware on the defective line module is still usable. The exception might have been triggered by a hardware fault that the diagnostic services might discover using its testing mechanism. For the stateful line module switchover to work correctly, the current subscriber information must be made accessible to the newly configured primary line module. The states of subscriber sessions that were on the failed line module are mirrored to the secondary line module, which has taken over as the primary controller.

Only those subscriber sessions whose operational status is up are guaranteed to be retained. Other subscriber connections that are in the process of fluctuating between up and down states might not be preserved. Such a technique of not retaining the subscriber sessions that are constantly alternating is adopted because of the possibility of occurrence of faults during transitional periods of states. If subscriber sessions in these transitional states are not preserved, the possibility of the same problem occurring on the secondary line module after the switchover is reduced.

Related Documentation

- Stateful Line Module Switchover Overview on page 68
- Stateful Line Module Switchover Platform Considerations on page 70
- System Operations When Stateful Line Module Switchover Is Enabled on page 74

Stateful Line Module Switchover Platform Considerations

Stateful SRP switchover is supported on all E120 and E320 routers that contain ES2 4G line modules installed with the ES2-ES1 Service IOA. See the *E120 and E320 Module Guide* for modules supported on the E120 and E320 Broadband Services Routers.

Table 15 on page 70 lists the line module, SRP module, and IOA slot combinations that support stateful switchover of line modules and stateful switchover for LNS sessions, when the router operates as an LNS device on one side of an L2TP tunnel.

Table 15: Module Configurations Supported for Stateful Switchover of LNS Sessions

Router Model	SRP and SFM Model	Number of L2TP tunnels and sessions	Number of Active and Standby ES2-ES1 Service IOAs	Downlink and Uplink LMs	Support for Stateful Switchover of LNS Sessions
E320	SRP-100	16,000 tunnels and 16,000 sessions	2 Service IOAs (1 active and 1 standby)	ES2 4G LM and GE-4 IOA	Supported
E320	SRP-100	16,000 tunnels and 32,000 sessions	4 Service IOAs (2 active and 2 standby)	ES2 4G LM and GE-4 IOA	Supported
E320	SRP-320	32,000 tunnels and 32,000 sessions	4 Service IOAs (2 active and 2 standby)	ES2 4G LM and GE-4 IOA	Supported
E120	SRP-320	16,000 tunnels and 16,000 sessions	2 Service IOAs (1 active and 1 standby)	ES2 10G LM and GE-8 IOA	Not supported
E120	SRP-320	32,000 tunnels and 32,000 sessions	4 Service IOAs (2 active and 2 standby)	ES2 10G LM and GE-8 IOA	Not supported
E120	SRP-320	16,000 tunnels and 16,000 sessions	2 Service IOAs (1 active and 1 standby)	ES2 10G LM and GE-8 IOA	Supported
E120	SRP-320	32,000 tunnels and 32,000 sessions	4 Service IOAs (2 active and 2 standby)	ES2 10G LM and GE-8 IOA	Supported

Related Documentation

- Stateful Line Module Switchover Overview on page 68
- System Operations When Stateful Line Module Switchover Is Enabled on page 74
- Replacement of Line Modules When Stateful Line Module Switchover Is Enabled on page 77
- Application Support for Stateful Line Module Switchover on page 79

Guidelines for Configuring Stateful Line Module Switchover

Keep the following points in mind when you configure stateful switchover for line modules:

- Line module high availability, similar to dual SRP stateful switchover, does not prevent the root cause of the reload or restart. This functionality is designed to return the system to the online, active state as soon as possible. The secondary line module takes the role of the primary module to preserve subscriber sessions in the up state with minimal subscriber data outage.
- The architecture of line modules supports switchover simultaneously across multiple 1:1 sets of line modules that participate in line module high availability.
- Line module high availability is available only on the operational image that runs on the interface controller (IC). The behavior of the forwarding controller (FC) image, and the IC boot and diagnostic images provide the same functionality as the behavior that existed before line module high availability support was implemented.
- Line module high availability in a 1:1 redundancy model is supported for ES2 4G LMs and Service IOAs on E120 and E320 routers. The architecture of line module high availability ensures that it does not depend on high availability for SRP modules to be enabled and operational for stateful line module switchover to work. Similarly, any modifications made to the dual SRP stateful switchover settings do not require or depend on line module high availability to be enabled and operational.
- Unified ISSU is supported on the primary line module in a high availability pair of line modules. The secondary line module is disabled during the unified ISSU operation and cold boots after the unified ISSU operation is complete. Line module high availability mode is active after the secondary line module is up, provided that line module HA configuration is enabled.
- Applications that are configured on the router ensure that their defined settings and memory requirements are handled on E120 and E320 routers. The primary and secondary line modules in a high availability pair are determined using the slot information specified using the **mode high-availability slot** command in Redundancy Configuration mode.
- Packets that are transmitted between the FC and IC and between the FC and system controller (SC) are not preserved during a stateful line module recovery.
- 1:N hot standby mode is not supported for stateful line module switchover. Automatic switchover of the serial connection to the line module that is designated as the primary module after switchover is also not supported. Similarly, cold switchover is not supported if line module high availability is not configured.
- Recovery of routers from double failures, such as simultaneous switchover of SRP and line modules, is not supported. Application-specific statistical details are not retained across a stateful line module switchover.
- Subscriber sessions that constantly move between up and down states are not maintained across a stateful switchover.
- If the line module that contains a downlink interface (connecting to the LAC device) reloads, owing to hardware or software failures, subscriber sessions are not maintained, even if the LM and Service IOA are HA-protected. Also, subscriber sessions are not retained if the line module that connects to the LAC device reloads, when the LM and Service IOA are part of a redundancy group.

- ES2 10G LMs cannot be used as downlink modules in an LNS device. These LMs cannot be used as access modules in a LNS device that contains a Service IOA that is HA-enabled.
- Certain statistics might be lost during the period of the stateful line module switchover. PPP and policy statistics are polled and collected every 10 minutes and sent to the standby line module. The statistics that were last collected before the switchover occurred are used as the baseline for statistics on the newly configured primary module. At a maximum, statistics for around 10 minutes might be lost. This scenario normally happens when polling is about to happen and the primary module switched over.
- A historical record of information about the forwarding and drop events and forwarding and drop rates on egress queues is not retained across a stateful line module switchover. The queue statistics for subscriber interfaces are calculated afresh after a stateful switchover of line modules.
- Sequence number checking for data packets received on all L2TP tunnels in the router is not maintained and supported during a stateful line module switchover. We recommend that you set up the router to ignore sequence numbers in data packets received on L2TP tunnels by entering the **l2tp ignore-receive-data-sequencing** command on an LNS device to prevent requests from a LAC device to enable insertion of sequence numbers into data packets.
- Some performance impact might occur when a new secondary module is provisioned or inserted, with the primary module containing maximum tunneled PPP sessions. In this case, data synchronization consumes a portion of the backplane bandwidth, which might have some impact on call setup rate (CSR) during this time. Under peak load conditions, it might take about 20 minutes for the system to become HA-active for Service IOAs.
- Removal of the Service IOA from the primary ES2 4G LM without powering it down does not trigger stateful line module switchover.
- The PPP application on ES2 4G LMs with Service IOA supports line module high availability. PPP session data is mirrored to the standby line module to attain high availability. The PPP application replicates the PPP sessions on the standby module and retains them across switchovers, in addition to accounting statistics. During line module switchover, the forwarding controller (FC) in the access module of the router that works as the LNS attempts to prevent timeouts of PPP sessions (due to the lack of PPP echo reply messages to the active subscribers) by sending echo response packets until the switchover is successfully completed.
- Policy manager is stateful line module switchover safe. Policy manager downloads the policy attachments from the SRP to the newly active line module after a switchover operation is detected. Policy statistics are preserved and made available across line module switchovers.
- The QoS application accomplishes the stateful line module switchover functionality by restoring the queues on subscriber interfaces in the newly active line module when the previously designated primary line module fails.
- During stateful line module switchover, the forwarding controller (FC) in the access module on the router functioning as the LNS device prevents timeouts of PPP sessions

owing to the absence of PPP echo reply messages in response to echo requests received from clients.

- Only two pairs of primary and secondary line modules can be configured on a single chassis for stateful switchover. As a result, only two line modules can be HA-safe. If high availability is activated, when the secondary module takes over as the primary module, existing subscribers are retained. If high availability is not activated, when the primary line module fails, the standby line module processes the regular router functions, but previously active subscriber sessions are not retained.
- Stateful line module switchover can be triggered when one of the following actions is performed on the primary line module, with high availability for line modules enabled on the router:
 - Disabling the module in the specified slot using the **slot disable** command
 - Rebooting a module in a selected slot on the router using the **reload slot** command
 - Performing a graceful switchover to the secondary line module using the **line-card switch** command
- If both the primary and secondary modules are cold booted (for example, when a chassis is cold started), and if the primary module does not become online for 8 minutes, the secondary module takes the role of the primary module. This behavior is similar to the line module redundancy mechanism.
- If high availability for line modules is active, the switchover is stateful. Subscribers are not disconnected and none of the existing client sessions are terminated or locked out during the line module switchover. A data traffic outage of about 2 minutes occurs, although subscribers are not disconnected. PPP echo requests from the subscribers are responded by the access module itself during the switchover period. This method works properly even if LAG interfaces are configured to connect to a LAC device.
- Information related to line module switchover is not forwarded to applications such as the AAA or RADIUS servers. These modules are not requested again for any accounting or authorization information for the same subscribers that were connected during the time of switchover.
- When the unified ISSU process is in progress, you cannot configure high availability for line modules if the initialization state of the unified ISSU operation has started. You must wait until the unified ISSU procedure is completed to enable high availability for line modules.
- Line module high availability does not interfere with the configurations made for unified ISSU and stateful SRP switchover functions. The secondary module in a line module high availability pair does not participate in the unified ISSU operation and is disabled during the upgrade process. The secondary module is cold started after the unified ISSU procedure is completed. However, the primary line module takes part in the unified ISSU process and undergoes a warm restart.
- PPP-based stacks (L2TP, PPP, and IP applications) for both IPv4 and IPv6 interfaces support stateful line module switchover.
- You can manually switch between the primary and secondary modules. While the secondary module attempts to take over as the primary module during a switchover,

if the secondary module fails to transition as the primary module within 5 minutes, the secondary module is cold booted.

- SNMP traps are generated after the switchover of the primary line module.
- Similar to dual SRP configuration and high availability of SRP modules, stateful line module recovery does not prevent the root cause that caused a router reload or stoppage of functioning. Stateful line module recovery enables the system to be returned to the fully functional state as soon as possible. If the conditions that caused the problem recur after a restart, an abrupt reload of the router might occur again. Stateful line module recovery minimizes forwarding impact on a restart to maximize customer uptime and causes the loss of packets during a restart to be limited to a small number of packets that are dropped in a timespan of a few seconds.

**Related
Documentation**

- System Operations When Stateful Line Module Switchover Is Enabled on page 74
- Application Support for Stateful Line Module Switchover on page 79
- Stateful Line Module Switchover Modes on page 83
- Stateful Line Module Switchover States on page 84
- Guidelines for Activating High Availability on page 87
- Activating High Availability on page 88
- Guidelines for Deactivating High Availability on page 89
- Deactivating High Availability on page 89

System Operations When Stateful Line Module Switchover Is Enabled

Line modules on E120 and E320 routers comprise three components—forwarding controller (FC), interface controller (IC), and the input/output adapter (IOA). The IC operational image runs in the IC section and the FC operational image runs in the FC section. Each line module can be connected to multiple IOAs, although the IOA does not contain an active control processor and has no operational image running on it. The IOA has a control path to the IC, which is used to configure and preset the hardware on the IOA. The IOA has a data path between the FC and the IOA external connections. After the hardware in the IOA is configured, it transfers data between the FC and the IOA external connections.

Line module high availability or stateful line module switchover behavior refers to providing this functionality of uninterrupted connectivity for subscribers by switching over to the operational image running on the IC that is active on the secondary line module of a pair of modules enabled for high availability. The applications on the secondary line module recover to their original states by reconstructing data from the preserved mirrored storage containers to a stable state. After the secondary line module is equipped to take over the load (services) of the primary line module, the switchover of subscriber traffic occurs on the secondary line module. The secondary line module begins to operate normally, although certain users might experience some subscriber data outage.

The design architecture used for E120 and E320 routers causes the packets that are designated for the SC to be first sent to the IC using a direct memory access (DMA) method. These packets are then forwarded to the SC over the internal 100 Mb Ethernet channel. Similarly, packets that are destined to be transmitted from the router are first sent to the IC over the Ethernet channel and then sent from the IC to the FC using a DMA operation. During the switchover period, until the secondary line module becomes operational, this communication channel from the IC is not active. The amount of time taken for the operational image on the secondary line module to start fully functioning can take up to several seconds. During this period of completion of the switchover process, applications on the SC handle this timeout. Applications that are impacted by this outage are routing protocols or protocols that are time-critical. For example, SRP-based TCP and UDP services are preserved across a switchover, which enables applications, such as Telnet, FTP, SSH, and SNMP, to operate seamlessly.

When the high availability functionality for line modules is active, subscriber sessions are maintained whenever a software or hardware fault is detected on the primary line module. A brief interruption occurs in the subscriber data traffic during the time that the secondary line module takes over as the primary line module.

Related Documentation

- Stateful Line Module Switchover Overview on page 68
- Guidelines for Configuring Stateful Line Module Switchover on page 70
- Stateful Line Module Switchover Platform Considerations on page 70
- Replacement of Line Modules When Stateful Line Module Switchover Is Enabled on page 77
- Application Support for Stateful Line Module Switchover on page 79
- Stateful Line Module Switchover Modes on page 83
- Stateful Line Module Switchover States on page 84
- Activating High Availability on page 88
- Deactivating High Availability on page 89

Stateful Line Module Configuration Scenarios

The line module high availability functionality is based on the stateful SRP switchover design architecture, with the implementation extended to two pairs of line modules to function as the primary and secondary module. This section describes the behavior of different system functions when line module high availability is configured on the router.

High Availability Configured and Enabled on the Line Module

If line module high availability is configured and the state is active, when a fault occurs on the primary line module, the primary line module performs a warm switchover to a secondary module. After the switchover, the secondary module starts operating as the new primary module. The previously configured primary module, after it becomes operational, takes over the role of the secondary module.

When high availability is enabled on a router, the secondary module takes over the role of the primary module, which causes normal system services and subscriber data traffic to continue without interruption after a switchover. The system controller (SC) on E120 and E320 routers is referred to as the SRP module and the main processor in a line module is called the interface controller (IC).

High Availability Configured and Disabled on the Line Module

If line module high availability is configured and the state is not active, when a fault occurs on the primary line module, the primary line module performs a cold switchover to a secondary module. After the switchover, the secondary module starts operating as the new primary module. The previously configured primary module, after it becomes operational, takes over the role of the secondary module. Subscribers are disconnected and need to log in again to establish their connections again. This behavior is similar to the functionality experienced during the line module redundancy operation.

High Availability Configured and the Switchover State Is Active or Disabled

Any failure on the secondary line module or a restart of the module causes high availability to move to the disabled state. The **show redundancy line-card** command displays the HA status on all line modules in the system. When a line module transitions from one state to another, log messages are seen on the SRP console and SNMP traps, if enabled, are sent.

All CLI commands that cause the line module to cold restart behave in the same manner, except for the method adopted to trigger the switchover action. For example, the **reload slot** and **slot disable** commands that reload a primary line module, causing the secondary module to take over as the primary. However, the **slot erase** command clears the configurations on the line module that is fitted in that slot. If you erase the slot configuration for a primary line module, both the primary and secondary line modules that participate in HA are cold booted. If you erase the slot configuration for the secondary line module, only that module is cold booted and the primary module is not cold booted.

If you enter the **reload slot** command, after booting up, line modules start operating in HA mode. With the **slot disable** command, HA is disabled until you reenables the slot.

Rebooting of the System When Line Module High Availability Is Configured

The line modules undergo a cold start, when the router is rebooted, and the secondary line module is held in a state in which it is not online. The primary line module reaches the online state. If the primary line module fails to come up online, within the specified timeout value (of less than 8.5 minutes), the secondary line module takes over as the primary module and HA remains in the disabled state.

Stateful SRP Switchover

During a stateful SRP switchover, a window of time occurs when the communication between interface controllers (IC) is disrupted owing to the switchover to new SRP module, which requires the Ethernet switch to relearn the MAC addresses and the interchassis communication (ICC) sessions to be reestablished. The system infrastructure ensures this task of relearning of details is transparent to the applications. Any notification sent by the applications on the IC-IC communication is either buffered until the

communication is reestablished. Otherwise, the Ethernet switch on the standby SRP module that has become active learns the MAC addresses in standby mode without interrupting the IC-IC communication.

Line Module Redundancy

Line module redundancy and line module high availability are mutually exclusive features. You cannot configure the line modules in a redundancy group to operate in HA mode. Because both the mechanisms are mutually exclusive, if a module is a member of a redundancy group, warm switchover is not supported. Similarly, if line modules are configured in a high availability pair, they cannot be members of a redundancy group and the spare module does not take over as the primary.

Unified ISSU

Unified ISSU can continue, if the configured secondary line module takes over as the primary line module. The secondary line module is disabled during the upgrade phase of the unified ISSU operation and cold boots after the unified ISSU operation is complete. The disabled line module during unified ISSU is cold booted after the unified ISSU operation is complete. Only the primary line module can participate in a unified ISSU operation.

Related Documentation

- Replacement of Line Modules When Stateful Line Module Switchover Is Enabled on page 77
- Application Support for Stateful Line Module Switchover on page 79
- Stateful Line Module Switchover Modes on page 83
- Stateful Line Module Switchover States on page 84
- slot erase
- slot disable
- reload slot

Replacement of Line Modules When Stateful Line Module Switchover Is Enabled

You can use the **show version** command to view the status of the SRP modules and line modules, as well as the running and armed releases. The state field in the output of this command indicates whether the slot that contains the redundant line module indicates standby or is restarting. Any CLI or SNMP command that the system attempts to process at the time of restart or recovery of the line module might fail. If you configured the ha logging event category (using the **log severity severityValue ha** command), log messages are displayed to specify that the line module has warm restarted or cold started. Depending on the nature of failures, the line modules that participate in HA take the following actions:

Reloading the Primary Line Module in Response to Failures

When a software fault occurs on the primary line module or if you enter the **slot disable** or **reload slot** commands for the slot in which the primary line module is installed, the

secondary line module takes over by recovering all the applications running on the primary module to a stable state. After the secondary module takes over, traffic starts flowing through the secondary module. The FC and IOAs on the faulty line module are cold booted along with the IC. When you use the **slot disable** command for a slot that contains a line module, you disable only the line module; you do not disable the line module and IOAs associated with it.

Reloading the Secondary Line Module in Response to Failures

When a software fault occurs on the secondary line module or if you enter the **slot disable** or **reload slot** commands for the slot in which the secondary line module is installed, the secondary line module undergoes a cold boot. After the secondary module becomes operational, the reloaded module continues to function as the secondary line module. Any core dumps that are generated on the faulty line module are sent to the SRP module, which is the same behavior as the one that occurs when a line module that is not configured for HA resets. When you use the **slot disable** command for a slot that contains a line module, you disable only the line module; you do not disable the line module and IOAs associated with it.

Disabling the Primary and Secondary Line Module Slots

If you specify the **slot erase** command on the primary line module to delete the configuration of the module in the selected slot before you install a different type of module, both the primary and secondary line modules are cold booted and line module high availability is disabled after the modules become operational again. If you enter the **slot erase** command on the secondary line module, only that module is cold booted and line module high availability is disabled.

Reloading the Router When Line Modules Enabled for HA Are Installed

If you enter the **reload** command on the router to restart the device with the currently available configuration, the previously configured roles of the primary and secondary line modules are preserved. SRP modules reload all line modules and restart from saved configuration files. However, any failure of the primary line module to become operational after the prescribed timeout duration causes the secondary module to take over as the primary.

Removing IOAs Without Powering Down from Line Modules

Removal of an IOA from the primary line module without powering down the IOA does not trigger line module switchover. Removal of an IOA from the secondary line module does not cause the module to be cold started. In both such scenarios, we recommend that you perform a hot-swap of the IOA in a managed environment when high availability is configured on the line module pair.

Cold and Warm Switchovers of Line Modules In a High Availability Pair

Cold switchover of the line module results in the same behavior as the system operations that are witnessed with line module redundancy configured on a router. With both these features, all the existing subscriber sessions are lost during the switchover.

When a warm switchover of the line module in a high availability pair occurs, subscriber sessions are not lost during the switchover.

Related Documentation

- Application Support for Stateful Line Module Switchover on page 79
- Log Messages Generated for Stateful LM Switchover on page 91
- Monitoring the Redundancy Status of Line Modules in a Specific Slot on page 95
- Monitoring the Redundancy History of Line Modules in a Specific Slot on page 97
- Interoperation of Redundancy and Stateful Switchover for Line Modules on page 15
- mode
- line-card switch
- show redundancy line-card
- show redundancy history

Application Support for Stateful Line Module Switchover

Applications are either supported or unsupported by stateful line module switchover.

- **Supported**—You can configure supported applications without having any adverse impact to stateful line module 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 line module 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.

The sections that follow describe the working behavior of applications that support stateful line module switchover.



NOTE: Only the applications discussed in the sections that follow are compatible with stateful line module switchover.

Policy Management

Because the policy application in the line module does not contain the complete state of all the policy definitions in mirrored containers, the SRP module is used to download the policy definitions and attachments to the newly active line module when a stateful switchover occurs. The policy application sends multiple policy attachment requests from the SRP module to the line module in a single notify operation and in a bulk manner, instead of one policy attachment request in each notify event. This method of transferring

policy attachment requests in bulk reduces the time to download all the attachments to the newly active line module.

QoS

QoS configuration is maintained in each line module and these settings are mirrored to the standby line module. During a stateful line module switchover, the QoS agent in the line module restores the configuration in the newly active line module. The QoS agent clients (such as IP and Ethernet) bind and register to the QoS agent before they replay the interfaces for creating QoS attachments. The QoS agents ensure that the queues are reestablished appropriately for the interfaces.

Connection Manager and Queue Manager

The queue manager resides on the SRP and the queue manager agents are present on all the line modules. When the primary line module resets, the spare module takes over the usage of the redundancy database. The queue manager identifies a connection based on the queue ID, the connection manager uses the stream ID to recognize a connection, the forwarding controller uses the stream ID, similar to the connection manager, to determine a connection. For example, when slot 2 communicates with slot 1, the queue manager identifies this connection as QID1. Similarly, when slot 3 communicates with slot 2, this link is labelled as QID2.

The connection manager uses SID1 to denote the connection from any slot with slot 2 and SID2 to signify the link from any slot with slot 3. The slot 2 address is specified as 2a2, where '2' refers to the logical slot, 'a' indicates the active state of the slot, and '2' represents the physical slot. When slot 0 takes over slot 2, the slot that is taken over is identified as 2a0. On reception of the controller up event on the SRP module for the spare line module, the queue manager initiates a request to the connection manager to create a fresh connection for the address 2a0. The connection manager logically labels the stream ID that refers to slot 2 to be down and creates a new stream ID to communicate with slot 0. The forwarding controller database that possesses a mapping of the slot ID, stream ID, and traffic class is updated accordingly to replace any streams that earlier pointed to slot 2 to start referring to slot 0. The queue manager agent running on the line modules handle the forwarding controller updates.

PPP

The PPP application on the line module contains the basic protocol, timers, and state machines in a running state. All the dynamic session data collected from protocol negotiations is present in the mirrored storage containers on the line module. For stateful line module switchover, all the mirrored storage data is saved on the standby module, replicating the session on the standby module. After the switchover takes place, the application initialization process on the standby module reconstructs the mirrored data and brings up the sessions to the established state (operational status is up). Some of the sessions that are still in the process of being created (alternating between the up and down operational states) are not retained during the switchover. This behavior of not preserving sessions that are not established is similar to the characteristic followed during unified ISSU, where sessions that are not completely created retry after the newly configured primary line module is available.

The total time required for the standby module to become active is dependent on the size of the configuration parameters. On a normal basis, it takes about 2-3 minutes for the new primary module to become active, in which case, clients running small intervals of keepalives expire. This system of expiry of keepalives poses a limitation on the stateful switchover model. This limitation is similar to the restriction seen during the upgrade phase of the unified ISSU process in which traffic forwarding is interrupted for a brief period. To work around this restriction, echo requests for the sessions that terminate on the failed line module are redirected to a different hardware. For failures on tunnel server modules (ES2 4G LMs with Service IOA), the access module handles such problems.

L2TP

L2TP configuration and operation data are maintained in the line module and this information is mirrored to the standby module. After the switchover of the primary tunnel server module to the secondary module occurs, the L2TP application on the line module restores the configuration and operation data to the newly active primary module. This mechanism is similar to the warm start procedure during unified ISSU. The L2TP application on the SRP module handles the line module events related to the primary and secondary modules.

Forwarding Controller

When a stateful line module switchover occurs, the forwarding controller (FC) tables that refer to the failed line module are updated with stream IDs that map to the line module (ES2 4G LM with Service IOA) that has taken over the role of the primary module. FC tables use a combination of slot ID, stream ID, and key hash table. The modifications to the FC tables enables packets to be sent to the newly functioning primary module after the switchover is complete.

During the stateful line module switchover, PPP subscriber sessions on an LNS device in an L2TP tunnel might be terminated due to the lack of PPP keepalive responses from the LNS device. To prevent the termination of subscriber sessions, the access module in the LNS device handles the PPP echo requests from all active subscriber sessions (on behalf of the failed line module) and responds with valid PPP echo reply messages. After a successful switchover, the access module in the LNS stops responding to the PPP echo request messages.

When the access module in the LNS receives an event from the application, such as PPP, to denote a failure with the primary line module, the access module starts processing the PPP echo requests that are destined for the LNS. The access module in the LNS concludes the handling of PPP echo requests after it receives a notification that the switchover is complete.

The following configuration events also take place during a stateful switchover on tunnel server modules that are installed on E120 and E320 routers that operate as LNS devices in an L2TP tunnel:

- All possible next hop attributes, which signify the IP address of the node that is closer to the advertised prefix (such as MPLS and ATM sessions), at the LNS are supported.
- PPP keepalive messages are not considered for the session statistics calculated during stateful switchover.

- Only the PPP echo request messages received on the L2TP tunnels or sessions that terminate at the LNS are handled by the access module during switchover. The FC in the access module in the LNS device does not send or generate any PPP echo request messages on its own.
- Sequence number checking for data packets received on all L2TP tunnels in the router and L2TP over IPSec to configure secured transport connections are not supported during a stateful line module switchover.
- During the switchover, when the access module that handles the echo request messages on the LNS fails (stops responding or traffic stops flowing), the PPP subscriber sessions that wait for echo response messages from the LNS terminate owing to the absence of a response.
- If line module redundancy is enabled and a switchover is being performed on an access module in a LNS device, and if a stateful line module switchover also commences at the same time, echo replies are not sent from the access module in the LNS. The PPP subscriber sessions that expect the echo response messages from the LNS during the switchover are terminated owing to the absence of an echo response.
- During a stateful line module switchover, if the secondary tunnel server module (ES2 4G LM with Service IOA and configured on a router that acts as the LNS) encounters a fault, the access module stops responding to PPP echo request messages after it receives the notification from the SRP module or the PPP application.

Mirroring Subsystem

The mirroring application is used to synchronize the configuration information available on the line modules. The mirroring state machine resides on both the primary and secondary line modules. The mirroring functionality uses interchassis communication (ICC) sessions to coordinate between line modules. Mirroring is supported for the volatile memory present on the line modules. After an initial bulk synchronization of storage data from the primary line module to the secondary line module occurs, any subsequent data is mirrored as and when transactions are posted. When a stateful switchover occurs, applications recover to the steady state by restoring the configuration data from the mirrored containers.

State machine-dependent applications, such as PPP, L2TP, and QoS applications, contain a dummy forwarding controller database that is populated on the access line module (receives traffic from low-speed circuits and routes them to uplink modules). This dummy database enables responses to be sent from the access line module to the keepalives that it receives until the switchover completes. This method of sending responses to hello packets ensures minimal data outage during the switchover of line modules. After the stateful switchover, the stateful applications start their regular processing by reestablishing their containers and perform a synchronization with the SRP module for dynamic data.

Unified ISSU

A unified ISSU operation proceeds properly if the configured secondary line module had taken over as the newly active primary line module. When you enter the **issu start** command to begin the upgrade phase of the unified ISSU process, the secondary line

module is disabled. The disabled line module during unified ISSU is cold booted after the unified ISSU operation is complete. Only the primary line module participates in the unified ISSU operation.

ICCP

Interchassis Communication Protocol (ICCP) is used to establish communication sessions between line modules that are configured for stateful switchover (configured in the high availability pair). Controller events are generated for existing sessions on the line modules with a notification about the session establishment and session teardown. The applications that are running on the SRP module with ICC sessions formed between the SRP and line modules are notified with the controller events after a stateful line module switchover occurs.

The line module high availability manager resides on the SRP module to enable the stateful switchover from a failed primary module to the secondary module in a high availability pair of devices. The high availability manager interacts with its peer agent on the line modules using ICC session and control bus. After the modules in a high availability pair become operational in primary and secondary modes, the high availability manager notifies interchassis controller (ICC) to enable ICC communication between the line modules.

Related Documentation

- Stateful Line Module Switchover Modes on page 83
- Stateful Line Module Switchover States on page 84
- Activating High Availability on page 88
- Deactivating High Availability on page 89

Stateful Line Module Switchover Modes

The line modules can operate in one of two redundancy modes—stateless switchover and high availability.

Stateless Switchover Mode

Stateless switchover is the default behavior mode for E120 and E320 routers that contain redundant line modules. This mode is available on line modules that you configured to function in a redundancy group and on which the high availability behavior is not enabled using the **mode high-availability slot** command. During switchover, the line, circuit, and IP interfaces on 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.

High Availability Mode

The high availability mode uses an initial bulk file transfer and subsequent, transaction-based mirroring to ensure rapid line module recovery after a switchover. This process is referred to as stateful line module switchover.

High availability mode keeps state and dynamic configuration data from the SRP memory synchronized between the primary and standby line modules.

When stateful line module switchover is enabled, user sessions are unaffected and forwarding of data is briefly disrupted, and the newly active line module continues from the point of switchover.

In high availability mode:

- The contents of the volatile storage in the primary and standby line modules remain synchronized using the mirroring methodology.

If a switchover occurs:

- The standby line 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.

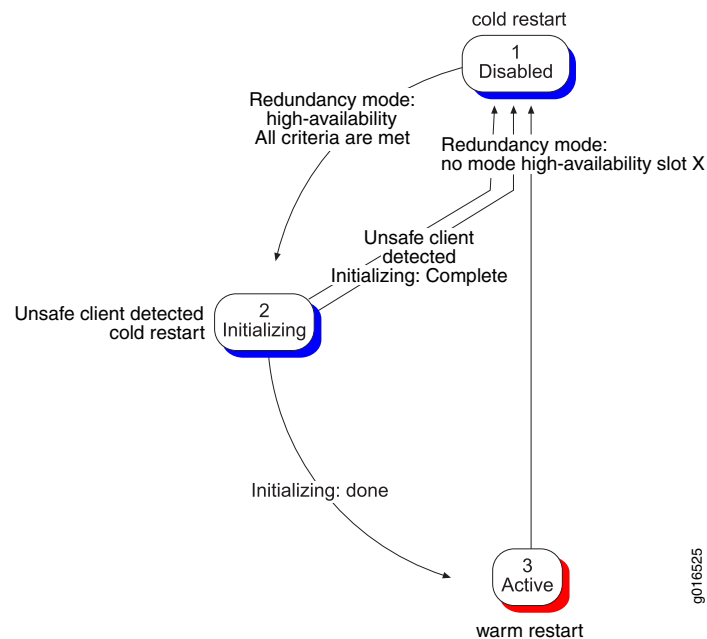
**Related
Documentation**

- Stateful Line Module Switchover States on page 84
- mode
- line-card switch
- show redundancy history
- show redundancy line-card

Stateful Line Module Switchover States

The line module progresses through various high availability states. These states are illustrated in Figure 5 on page 85.

Figure 5: Stateful Line Module Switchover States



Disabled State

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

The router enters this state when you power up the router or when the router warm-restarts from a stateful line module 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 line module hardware supports high availability.
- Running configuration allows high availability to operate (that is, no unsupported applications are configured).
- Standby line module hardware supports high availability.
- Standby line module is online and capable of mirroring.
- Standby line module is running the same release as the active line module.

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 undergoes a cold-switchover.

The behavior of the router depends on which HA state the application is in when it shifts to a disabled state:

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



NOTE: You can use the `show redundancy line-card` command to display the name of any unsupported applications that are configured.

While in the disabled state, the system operates as if it were configured for redundancy functionality.

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

Initializing State

After the line module transitions into the initializing state, bulk synchronization of the memory occurs. This activity consists of mirroring of appropriate state and dynamic configuration information from the active line module (memory) to the standby line module (memory). 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 disabled 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 undergoes a cold-switchover.
- After initialization is completed, the system enters the active state.

Active State

During the active state, the data that was synchronized from the active line module to the standby line 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 line module, and post the transactions.
2. Following the updates, the active line module sends the changes to the standby line module.

3. The standby line module replays the updates (in the order in which they were committed on the active line module) and makes the appropriate changes for each changed application.

During the active state:

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

**Related
Documentation**

- Stateful Line Module Switchover Modes on page 83
- Application Support for Stateful Line Module Switchover on page 79
- mode
- line-card switch
- show redundancy history
- show redundancy line-card

Guidelines for Activating High Availability

Before you activate high availability on the line modules, you must be aware of any high availability–related changes to line module management commands. For information on high availability–related changes to SRP modules, see “Managing Stateful SRP Switchover” on page 35.

You activate high availability (stateful line module switchover) by launching Redundancy Configuration mode and issuing the **mode high-availability slot** command. The **high-availability slot** keyword enables high availability mode for stateful line module switchover. In this mode, the router uses mirroring to keep the configuration and state of the standby line module coordinated with the configuration and state of the active line module.

When activating high availability, keep the following in mind:

- In an E120 router and an E320 router that supports stateful line module switchover, both line 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 unsafe configuration on the active and standby line modules, the system reverts to its default mode (redundancy configuration).
- When active, the router configuration files are mirrored from the active line module to the standby line module.

- Related Documentation**
- Activating High Availability on page 88
 - Deactivating High Availability on page 89
 - Guidelines for Deactivating High Availability on page 89
 - mode
 - line-card switch
 - show redundancy history
 - show redundancy line-card

Activating High Availability

The line module can operate in one of the two redundancy modes—stateless switchover and high availability. When you activate high availability, the router uses mirroring to keep the configuration and state of the standby line module coordinated with the configuration and state of the active line module.

To activate high availability:

1. From Global Configuration mode, launch Redundancy Configuration mode.
host1(config)#redundancy
2. In Redundancy Configuration mode, configure high availability as the redundancy mode and specify the line module pair (primary and secondary) for stateful switchover.

host1(config-redundancy)#mode high-availability 11 16

In this example, the line module that resides in slot 11 is assigned as the primary module and the line module installed in slot 16 is assigned as the standby or secondary module.

A warning message is displayed prompting you to confirm (enter **y** for yes, **n** for no) whether you want to proceed with the reloading of the secondary line module in the high availability pair. If you enter yes, the line module pair is set for stateful switchover after a reload operation is completed on the secondary line module. If you enter no, the high availability setting is not saved on the line module pair.



NOTE: You can specify only two pairs of line modules for high availability operation. If you attempt to configure high availability for more than two pairs of line modules, an error message states that the maximum number of line-module high availability pairs are already configured on the router and the setting is not saved.

- Related Documentation**
- Guidelines for Activating High Availability on page 87
 - Stateful Line Module Switchover Modes on page 83
 - Application Support for Stateful Line Module Switchover on page 79
 - mode

- line-card switch
- show redundancy history
- show redundancy line-card

Guidelines for Deactivating High Availability

You can disable high availability for line modules by launching Redundancy Configuration mode and issuing the **no mode high-availability slot** command. In the redundancy mode, the spare line module takes control of the IOA associated with the failed line module in the redundancy group. During switchover, the line, circuit, and IP interfaces on one or more IOAs appear to go down temporarily. This is the default behavior mode for E120 and E320 routers that contain redundant line modules.

Because this mode requires the router to reload the interface configuration and the routing table from the SRP module, instead of transaction-based mirroring, when the active line module switches to the standby line module, a disruption occurs in handling user sessions and traffic forwarding.

Related Documentation

- Deactivating High Availability on page 89
- Stateful Line Module Switchover Modes on page 83
- Application Support for Stateful Line Module Switchover on page 79
- mode
- line-card switch
- show redundancy history
- show redundancy line-card

Deactivating High Availability

When you disable stateful line module switchover, both the primary and standby line modules are shifted to the online state. Both line module redundancy and stateful switchover of line modules are disabled. If the router uses redundant line modules, after you disable high availability, the router uses line module redundancy mode, which is the default behavior.

To disable high availability support:

1. From Global Configuration mode, launch Redundancy Configuration mode.
`host1(config)#redundancy`
2. In Redundancy Configuration mode, disable high availability on the line module pair previously configured for stateful switchover
`host1(config-redundancy)#no mode high availability slot 11 16`

In this example, high availability behavior configured for the primary line module that resides in slot 11 and the line secondary module installed in slot 16 is removed.

In this scenario, if slot 11 is in online state and slot 16 is in standby state, a warning message is displayed prompting you to confirm (enter **y** for yes, **n** for no) whether you want to proceed with the reloading of the line module in slot 16 and whether you want the primary and secondary modules to be removed from the high availability pair. If you enter yes, the line module pair is disabled for stateful switchover after a reload operation is completed on slot 16. If you enter no, the high availability setting continues to be activated on the line module pair.

Alternatively, if the secondary line module in slot 16 has taken over the role of the primary line module in slot 11, and if you disable high availability on the line module pair, a warning message is displayed prompting you to confirm (enter **y** for yes, **n** for no) whether you want to proceed with the reloading of both the line modules in slot 16 and slot 11 and whether you want the modules to be removed from the high availability pair. If you enter yes, the line module pair is disabled for stateful switchover after a reload operation is completed on slot 16 and slot 11. If you enter no, the high availability setting continues to be activated on the line module pair.

Related Documentation

- Guidelines for Deactivating High Availability on page 89
- Stateful Line Module Switchover Modes on page 83
- Application Support for Stateful Line Module Switchover on page 79
- mode
- line-card switch
- show redundancy history
- show redundancy line-card

Switching Over from a Primary Line Module to Secondary Line Module

You can enable a manual switchover from the currently functioning primary module to the secondary module, even when the primary module has not encountered a fault, using the **line-card switch** command in Privileged Exec mode.

To perform a stateful switchover from the currently active primary module to the secondary module when high availability for line modules is enabled:

- Specify the slot in which the line module is located.

host1#line-card switch 10



NOTE: You can perform a switchover only to ES2 4G LMs installed with the ES2-ES1 Service IOA on E120 and E320 routers. After you enter this command, the line module configured as the secondary starts functioning as the primary module. The previously configured primary module, after it becomes operational again, takes over the role of the secondary module.

The high availability operation on each line module is controlled by a state machine. In this example, the line module in slot 10 takes over as the primary module.

- Related Documentation**
- Stateful Line Module Switchover Modes on page 83
 - Application Support for Stateful Line Module Switchover on page 79
 - mode
 - line-card switch
 - show redundancy history
 - show redundancy line-card

Log Messages Generated for Stateful LM Switchover

When you configure high availability for line modules and SRP modules, the modules that support stateful switchover record log messages and SNMP traps when they transition to different switchover states, such as disabled, initializing, and active.

In releases of JunosE Software that did not support high availability for line modules, the logging messages for the ha event log category displayed information about high availability functions that were specific to SRP switchover operations. Although the log messages were associated with SRP switchover, the high availability details in the recorded messages did not explicitly indicate that the information was relevant to only SRP modules.

In JunosE releases that support high availability for line modules, the existing messages for the ha log event category have been enhanced to uniquely distinguish the changes to switchover states for line modules and SRP modules. The following examples describe logging messages generated when SRP modules and line modules transition from one state to the other during the switchover process.

Log Messages Displayed During the Transition from Disabled State to Active State

In the disabled state, which is the default state for high availability mode, the router uses redundancy mode for line modules (if a standby line module is configured for switchover purposes). If all criteria are met for the system to shift from the disabled state to begin bulk synchronization of the memory, high availability mode transitions to the initialization state. After initialization is completed, the system enters the active state. The following log messages are saved for SRP modules and line modules when the system moves from the disabled state to the active state. You must configure the high availability log event category using the **logging severity ha** command in Global Configuration mode to enable logs related to high availability operations to be recorded.

```
NOTICE 03/08/2005 10:34:35 ha: High Availability on the SRP is now active
NOTICE 03/08/2005 10:34:35 ha: High Availability on Slots X, Y is now active
```

Log Messages Displayed During the Transition from Active State to Pending or Disabled State

When an unsupported application is configured after HA is active on the line module, the application is regarded as unsafe for HA. The state machine on the line module transitions

to the disabled state. For SRP modules, the system transitions to the pending state. The following sample log messages are generated for SRP modules when the system transitions from the active state to the pending state, and for line modules when the system transitions from the active state to the disabled state:

WARNING 03/08/2005 11:49:21 ha: High Availability on the SRP is pending due to the creation of an unsafe configuration

WARNING 03/08/2005 11:49:21 ha: High Availability on slots X, Y is disabled due to the creation of an unsafe configuration

NOTE: At this point, the creation of an unsafe configuration on the SRP might not cause high availability on the line modules to change to the disabled state and vice-versa.

Log Messages Displayed During the Transition from Pending or Disabled State to Active State

The system transitions to the pending state if an unsupported application is configured for SRP modules. For line modules, when an unsupported application is configured, the system transitions to the disabled state. After the configuration of the unsupported application is removed, the system moves from the pending or disabled state to the active state. The following sample log messages are generated for SRP modules that are configured for high availability when the system transfers from the pending state to the active state, and for line modules when the system transfers from the disabled state to the active state:

WARNING 03/08/2005 11:54:10 ha: High Availability on the SRP is active due to removal of an unsafe configuration

WARNING 03/08/2005 11:54:10 ha: High Availability on cards in slot X, Y is active due to removal of an unsafe configuration

Log Messages Displayed During the Transition from Active or Pending State to Disabled State

From the active state, the system changes to the pending state if an unsupported application is configured for SRP modules. From the active state, the system changes to the disabled state for line modules if an unsupported application is configured or if any other criterion is not satisfied for remaining in the active state. The following log messages are generated for SRP modules and line modules that configured for high availability when the system transfers from the pending state or the active state to the disabled state:

WARNING 03/09/2005 12:41:40 ha: High Availability on the SRP is disabled. Please check redundancy configuration and status

WARNING 03/09/2005 12:41:40 ha: High Availability on the cards in slot X, Y is disabled. Please check redundancy configuration and status

Log Messages Displayed for Stateful SRP and Line Module Switchover When HA Is Enabled

The following log messages are displayed for SRP modules and line modules when stateful switchover is enabled for those modules.

ERROR 12/15/2003 23:34:23 os (): Standby SRP has assumed the role of Active SRP (warm restart)

ERROR 12/15/2003 23:34:23 os (): Secondary line card in slot X has assumed the role of Primary line card in slot Y (warm restart)

Log Messages Displayed for Stateful SRP and Line Module Switchover When HA Is Disabled

The following log messages are displayed for SRP modules and line modules when stateful switchover is disabled for those modules. Any failure or restart of the secondary line module causes high availability to move to the disabled state.

CRITICAL 12/16/2003 00:03:39 os (): Standby SRP has assumed the role of Active SRP (cold restart)

WARNING 12/15/2003 23:34:23 os (): Line card in slot X has cold-restarted.

Related Documentation

- Stateful Line Module Switchover Modes on page 83
- Stateful Line Module Switchover States on page 84

Preservation of Statistics During Stateful Line Module Switchover

Statistical details associated with PPP and policy manager applications are retained across a stateful line module switchover. All other counters for settings and parameters connected with other applications running on the router are reset to zero after the switchover occurs to the new primary module. The following sections describe details on the statistics polling process and the attributes that are maintained after a stateful line module switchover:

PPP Accounting Statistics

For each PPP subscriber, accounting statistics are polled and collected every 10 minutes. These statistics are synchronized with the secondary line module. The statistical details collected just before the switchover represent the baseline for statistics on the newly configured primary module. At a maximum, statistics that accumulate for about 10 minutes might be lost and not synchronized with the new primary module. This scenario occurs when polling was to happen and the primary module switched over to the secondary module, just before polling commenced. The JunosE software begins the collection of accounting statistics for terminated PPP sessions following, but not including, authentication acknowledgement from the E120 and E320 router. The acknowledgment is either a CHAP success or PAP acknowledgement packet. The following PPP statistics are retained across the stateful line module switchover:

- Number of IPv4 octets received from the interface (Acct-Input-Octets, RADIUS attribute 42)
- Number of IPv4 octets transmitted to the interface (Acct-Output-Octets, RADIUS attribute 43)
- Number of IPv4 packets received from the interface (Acct-Input-Packets, RADIUS attribute 47)
- Number of IPv4 packets transmitted to the interface (Acct-Output-Packets, RADIUS attribute 48)
- Number of IPv6 octets received from the interface (Ippv6-Acct-Input-Octets, RADIUS attribute 26-151)

- Number of IPv6 octets transmitted to the interface (Ipv6-Acct-Output-Octets, RADIUS attribute 26-152)
- Number of IPv6 packets received from the interface (Ipv6-Acct-Input-Packets, RADIUS attribute 26-153)
- Number of IPv6 packets transmitted to the interface (Ipv6-Acct-Output-Packets, RADIUS attribute 26-154)

Policy Statistics

Similar to PPP accounting statistics, aggregates of policy parameters are polled and collected every 10 minutes. These statistics are synchronized with the secondary line module. The statistical details collected just before the switchover represent the baseline for statistics on the newly configured primary module. At a maximum, statistics that accumulate for about 10 minutes might be lost and not synchronized with the new primary module. This scenario occurs when polling was to happen and the primary module switched over to the secondary module, just before polling commenced. The counters for all policy-managed packets and octets that are configured for simple classifier, rate-limit, and parent-group actions are retained across stateful line module switchover.

Related Documentation

- Application Support for Stateful Line Module Switchover on page 79

Performance Impact and Scalability Considerations

Some performance impact occurs when a new secondary line module is provisioned or inserted, with the primary module containing the maximum number of tunneled PPP sessions. The synchronization of data occupies some backplane bandwidth, which might slightly impact the call setup rate during the time the secondary module is plugged in to the slot.

The maximum number of ingress or egress policy statistic entries that can be configured is 512,000 entries. Each statistic entry is a combination of an 8-byte packet counter and an 8-byte octet counter. The configuration of this maximum number of policy statistics causes the collection of 8 MB of data (512,000 entries * 16 bytes = 8,192,000 bytes).



TIP: Because the collection of statistics and mirroring of data is processor-intensive, we recommend that you do not exceed 2 MB of statistics data when you configure stateful line module switchover. The preservation of 2 MB of statistical information maps to an average of 8 statistics entries per subscriber for a fully scaled subscriber environment (16,000 subscribers * 8 statistics per subscriber * 16 bytes per entry = 2,048,000 bytes). You can maintain the limit on 2MB of statistics by reducing the number of subscribers and increasing the statistics per subscriber, or by increasing the number of subscribers and reducing the statistics per subscriber.

Although the applications that support stateful line module switchover are enhanced in their infrastructure and design capabilities, performance of the line modules is not

impacted by the improvements made to the compatible protocols and applications. In a scaled topology, simultaneous switchover of all the line modules in a chassis is supported. The operations of line modules that do not support stateful switchover is also not impacted in any manner, owing to the changes made to HA-safe applications.

- Related Documentation**
- Application Support for Stateful Line Module Switchover on page 79
 - System Operations When Stateful Line Module Switchover Is Enabled on page 74

Use of Status LEDs to Monitor the High Availability States of Line Modules

You can determine the redundancy state of line modules and SRP modules by examining their status LEDs. The LED on the line module denotes the state of the line module. In addition, if you issue the **show version** command, the state field for the slot that contains the redundant line module indicates standby. The following are the functions of the LEDs for the different states of the line modules in a high availability pair:

- OK LED—This LED is turned off while the secondary line module is switching over
- ONLINE LED—There is no difference in the method of illumination of this LED, regardless of whether line module high availability is enabled.
- REDUNDANT LED—This LED is turned on for the secondary line module and is not illuminated for the primary line module.

- Related Documentation**
- Monitoring the Redundancy Status of Line Modules in a Specific Slot on page 95
 - Monitoring the Redundancy History of Line Modules in a Specific Slot on page 97

Monitoring the Redundancy Status of Line Modules in a Specific Slot

Purpose Display the supported redundancy modes as well as other status relating to high availability for a line module installed in a particular slot.

Action To display the redundancy information of a line module in a particular slot.

```
host1#show redundancy line-card slot 3
Line Card
```

```
-----
```

```
automatic reverting is off
```

slot	High Availability State	last activation type	hardware role	lockout config	backed up by slot	sparing for slot	revert at
----	-----	-----	-----	-----	-----	-----	-----
3	disabled	cold	---	---	---	---	---

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

slot	criteria	met
---	-----	---
3	Running configuration is safe for High Availability?	No

Clients with Unsafe Configurations for High Availability

client	mode	configuration
Global Ipv6	unsupported	unsafe

Meaning Table 16 on page 96 lists the **show redundancy line-card slot slotNum** command output fields.

Table 16: show redundancy line-card slot slotNum Output Fields

Field Name	Field Description
Line Card	Redundancy and configuration details for the line module in the specified slot
High Availability State	State of the high availability mode of the line module in the specified slot (disabled or active)
last activation type	Last type of activation that occurred on this router (that is, the method by which the line module last booted [cold-start or warm-start])
automatic reverting	Whether the router is enabled to revert from all spare line modules to the associated primary line modules automatically: on or off
slot	Slot number in which the line module resides
hardware role	Function of the line module: primary or spare.
lockout config	Status of redundancy on the line module: <ul style="list-style-type: none"> --- —Line module redundancy is not supported protected—Line module redundancy is enabled locked out—Line module redundancy is disabled
sparing for slot	Slot that contains the line module that is a spare for this primary line module
backed up by slot	Slot that contains the primary line module for which this module is a spare
revert at	Time at which you want the line module to revert
Criteria Preventing High Availability from being Active	Criteria preventing the router from being in the active state of high availability mode. For the router to be in the Active state, all criteria for this option must be "yes"
criterion	Name of the criterion that must be met for the system to enter the active state of the stateful line module switchover process. During the active state, the mirroring subsystem synchronizes the secondary line module with the primary line module by mirroring updates to mirrored volatile storage
slot	Slot number in which the line module resides

Table 16: show redundancy line-card slot slotNum Output Fields (*continued*)

Field Name	Field Description
met	Whether the criterion for high availability of the line-module pair is satisfied. Possible values: <ul style="list-style-type: none"> • Yes—The specific criterion has been met for the system to enter the active state • No—The specific criterion has not been met and the system transitions to the pending state
Clients with Unsafe Configurations for High Availability	Whether the client supports high availability and also the safety level of configuration
client	Name of the application that is running on the line module configured for high availability
mode	High availability status of the application running on the line module: supported or unsupported
configuration	Safety level of the configuration based on whether or not the client is supported or unsupported and in case of those unsupported, whether or not the client has been configured. For example, if an unsupported client has been configured on a router with high availability enabled, the configuration reads "unsafe"

- Related Documentation**
- Activating High Availability on page 88
 - Deactivating High Availability on page 89
 - mode
 - line-card switch
 - show redundancy history
 - show redundancy line-card

Monitoring the Redundancy History of Line Modules in a Specific Slot

Purpose Display information about dates, times, and the number of occurrences for starts and switchovers for a line module installed in a particular slot.

Action To display information about the number of occurrences for starts and switchovers for a particular line module:

```
host1#show redundancy history line-card slot 3
```

slot	system up time	last cold switchover	last warm switchover	Number of cold switchovers	Number of warm switchovers
----	-----	-----	-----	-----	-----
3	00:32:09	2009-06-10 17:56:15	2009-06-01 17:56:15	8	2

To display additional redundancy history information:

```
host1#show redundancy history line-card slot 3 detail
```

Line card

slot	system up time	last cold switchover	last warm switchover	Number of cold switchovers	Number of warm switchovers
3	00:32:09	2009-06-10 17:56:15	2009-06-01 17:56:15	8	2
LC activation time	type	slot	uptime		
2009-06-10 19:20:35	warm-switch	03	0 20:20:40		
2009-06-10 18:58:23	warm-switch	03	0 19:58:28		
2009-06-10 18:39:54	warm-switch	03	0 19:39:59		
2009-06-10 18:20:10	warm-switch	03	0 19:20:15		
2009-06-10 18:00:46	warm-switch	03	0 19:00:50		
2009-06-10 17:41:13	warm-switch	03	0 18:41:17		
2009-06-10 17:12:03	warm-switch	03	0 18:12:08		
2009-06-10 16:47:21	warm-switch	03	0 17:47:23		
2009-06-10 16:26:59	warm-switch	03	0 17:27:00		
2009-06-10 15:54:12	warm-switch	03	0 16:54:12		
2009-06-10 15:31:57	warm-switch	03	0 16:31:57		

Meaning Table 17 on page 98 lists the **show redundancy history line-card slot slotNum** command output fields.

Table 17: show redundancy history line-card slot slotNum Output Fields

Field Name	Field Description
Line card	Information about the switchover operations performed on the specific line module
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
Number of cold switchovers	Total number of cold switchovers the router has experienced
Number of warm switchovers	Total number of warm switchovers the router has experienced
LC activation time	Amount of time the line module has been active
type	Type of switchover. Possible options: <ul style="list-style-type: none"> cold-switch—When the router is in pending state and switchover occurs, the line module undergoes a cold-switch or cold re-start warm-switch—When the router is in active state and switchover occurs, the line module undergoes a warm-switch or warm restart
slot	Slot in which the line module resides
uptime	Amount of time the line module has been operational

- Related Documentation**
- Activating High Availability on page 88
 - Deactivating High Availability on page 89
 - mode
 - line-card switch
 - show redundancy history
 - show redundancy line-card

CHAPTER 5

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 JunosE 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 102
- Unified ISSU Platform Considerations on page 104
- Hardware and Software Requirements Before Beginning a Unified ISSU on page 104
- Unified ISSU Terms on page 106
- Unified ISSU References on page 107
- Unified ISSU Phases Overview on page 107
- Unified ISSU Initialization Phase Overview on page 108
- Unified ISSU Upgrade Phase Overview on page 110
- Unified ISSU Service Restoration Phase Overview on page 115
- Application Support for Unified ISSU on page 115
- Unexpected AAA Authentication and Authorization Behavior During Unified ISSU on page 124
- Unexpected ATM Behavior During Unified ISSU on page 124
- Unexpected DHCP Behavior During Unified ISSU on page 125
- Unexpected Denial-of-Service Protection Behavior During Unified ISSU on page 125
- Unexpected Ethernet Behavior During Unified ISSU on page 126
- Unexpected File Transfer Protocol Server Behavior During Unified ISSU on page 127
- IS-IS Effects on Graceful Restart and Network Stability During Unified ISSU on page 129
- Unexpected L2TP Failover of Established Tunnels During Unified ISSU on page 131
- OSPF Effects on Graceful Restart and Network Stability During Unified ISSU on page 132
- Unexpected Suspension of PIM During Unified ISSU on page 133
- Unexpected Suspension of Subscriber Login and Logouts During Unified ISSU on page 134
- Unexpected SONET and SDH Behavior During Unified ISSU on page 134

- Unexpected T3 Behavior During Unified ISSU on page 135
- Unavailability of TACACS+ Services During Unified ISSU on page 135
- Interruption in Traffic Forwarding for Layer 3 Routing Protocols During Unified ISSU on page 135
- Recommended Settings for Routing Protocol Timers During Unified ISSU on page 138
- Upgrading Router Software with Unified ISSU on page 139
- Halt of Unified ISSU During Initialization Phase Overview on page 142
- Halting Unified ISSU During Initialization Phase on page 142
- Halt of Unified ISSU During Upgrade Phase Overview on page 143
- Halting Unified ISSU During Upgrade Phase on page 143
- Monitoring the Status of the Router During Unified ISSU on page 144

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 JunosE 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 a unified ISSU-challenged application is configured, the unified in-service software upgrade cannot proceed. You must unconfigure the unified 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, gracefully.
- During the unified ISSU process, with a high availability line-module pair configured on a router, the primary line module supports unified ISSU. In such a scenario, the secondary line module is disabled when unified ISSU is performed and is cold booted after the unified ISSU procedure is complete. High availability mode is reactivated after the secondary line module comes back online, if the line module HA configuration continues to remain enabled on the secondary module. For more information about how the stateful line module switchover functionality behaves during a unified ISSU process, see "Managing Stateful Line Module Switchover" on page 67.

Related Documentation

- Unified ISSU Phases Overview on page 107
- Application Support for Unified ISSU on page 115
- Hardware and Software Requirements Before Beginning a Unified ISSU on page 104
- Upgrading Router Software with Unified ISSU on page 139

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.



NOTE: In JunosE Release 11.3.0, the design enhancements for reenabling CSR FPGA upgrades are implemented, which enables you to perform a stateless upgrade (non-unified ISSU method, with router-wide outage for users) from JunosE releases numbered lower than Release 11.3.0 to JunosE Release 11.3.0. However, you cannot run a unified ISSU procedure from JunosE releases that do not contain the design changes for CSR FPGA functionality (earlier than Release 11.3.0) to JunosE releases that contain the enhanced CSR FPGA design (Release 11.3.0 and later releases), if the router chassis contains ES2 10G LMs or REV-02 ES2 10G LMs.

In releases numbered lower than Release 11.3.0, you could not perform either a conventional upgrade or a unified ISSU process on ES2 10G LMs that contained the CSR FPGA upgraded image. You could only download the image from serial Programmable Read-Only Memory (PROM), which was a factory image.

Hardware and Software Requirements Before Beginning a Unified ISSU

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.

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 JunosE 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.

- 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 Stateful SRP Switchover” on page 35 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.

Related Documentation

- Application Support for Unified ISSU on page 115
- Application Support for Stateful SRP Switchover on page 42
- Activating High Availability on page 52
- `license issu`
- `show issu`

Unified ISSU Terms

Table 18 on page 106 defines terms relevant to module behavior during a unified in-service software upgrade.

Table 18: 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.

Table 18: Unified ISSU-Related Terms (*continued*)

Term	Meaning
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.
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 35.

For more information about SRP module redundancy, see “Managing Module Redundancy” on page 9.

Unified ISSU Phases Overview

The JunosE 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 “Hardware and Software Requirements Before Beginning a Unified ISSU” on page 104 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.

Related Documentation

- Unified ISSU Initialization Phase Overview on page 108
- Unified ISSU Upgrade Phase Overview on page 110
- Unified ISSU Service Restoration Phase Overview on page 115
- Halting Unified ISSU During Initialization Phase on page 142
- Halting Unified ISSU During Upgrade Phase on page 143

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`.

Related Documentation

- Unified ISSU Phases Overview on page 107
- Halt of Unified ISSU During Initialization Phase Overview on page 142
- Halting Unified ISSU During Initialization Phase on page 142
- `issu initialize`
- `issu start`
- `issu stop`
- `show issu`

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 19 on page 111 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 19: 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.

Table 19: Router Response to Undesirable Events During the Upgrade Phase (*continued*)

Event	Router Action
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 (post-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.

Verifications of Requirements

Because some time may have passed since unified ISSU verified the requirements for the upgrade during the initialization phase, unified ISSU reverifies 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 the applications running on the new release can interpret. 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 instructs 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.

Related Documentation

- Unified ISSU Phases Overview on page 107
- Halt of Unified ISSU During Upgrade Phase Overview on page 143
- Halting Unified ISSU During Upgrade Phase on page 143
- `issu start`
- `show issu`
- `halt`
- `reload`

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.

Related Documentation

- Unified ISSU Phases Overview on page 107

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 20 on page 116 indicates which applications support or do not support a unified in-service software upgrade, as well as limitations on their behavior.

Table 20: Application Support for Unified In-Service Software Upgrades

Application	Supported	Unsupported	Notes
Physical Layer Protocols			
DS1 (E120 and E320)	–	–	–
DS1 (ERX1440)	–	–	–
DS3	✓	–	–
HDLC	✓	–	–
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			
ARP	✓	–	<p>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.</p>

Table 20: Application Support for Unified In-Service Software Upgrades
(continued)

Application	Supported	Unsupported	Notes
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.
IPSec Transport (E120 and E320)	–	✓	E120 and E320 routers do not support IPSec.
IPSec Transport (ERX1440)	–	–	–

**Table 20: Application Support for Unified In-Service Software Upgrades
(continued)**

Application	Supported	Unsupported	Notes
IPSec Tunnels (E120 and E320)	–	✓	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.
DVMRP (E120 and E320)	✓	–	–
DVMRP (ERX1440)	–	–	–

Table 20: Application Support for Unified In-Service Software Upgrades
(continued)

Application	Supported	Unsupported	Notes
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 20: 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 20: 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	✓	–	–

**Table 20: Application Support for Unified In-Service Software Upgrades
(continued)**

Application	Supported	Unsupported	Notes
IOA hot swap	–	✓	–
J-Flow (IP flow statistics)	✓	–	–
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 20: 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	✓	–	–

Related Documentation

- Unexpected AAA Authentication and Authorization Behavior During Unified ISSU on page 124
- Unexpected ATM Behavior During Unified ISSU on page 124
- Unexpected DHCP Behavior During Unified ISSU on page 125
- Unexpected Denial-of-Service Protection Behavior During Unified ISSU on page 125
- Unexpected Ethernet Behavior During Unified ISSU on page 126
- Unexpected File Transfer Protocol Server Behavior During Unified ISSU on page 127
- IS-IS Effects on Graceful Restart and Network Stability During Unified ISSU on page 129
- Unexpected L2TP Failover of Established Tunnels During Unified ISSU on page 131
- Interruption in Traffic Forwarding for Layer 3 Routing Protocols During Unified ISSU on page 135
- OSPF Effects on Graceful Restart and Network Stability During Unified ISSU on page 132
- Unexpected Suspension of PIM During Unified ISSU on page 133
- Unexpected SONET and SDH Behavior During Unified ISSU on page 134

- Unexpected T3 Behavior During Unified ISSU on page 135
- Unexpected Suspension of Subscriber Login and Logouts During Unified ISSU on page 134
- Unavailability of TACACS+ Services During Unified ISSU on page 135
- Recommended Settings for Routing Protocol Timers During Unified ISSU on page 138

Unexpected AAA Authentication and Authorization Behavior During Unified ISSU

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+.

Related Documentation

- Application Support for Unified ISSU on page 115

Unexpected ATM Behavior During Unified ISSU

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.

Related Documentation

- Application Support for Unified ISSU on page 115

Unexpected DHCP Behavior During Unified ISSU

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

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.

Related Documentation

- Application Support for Unified ISSU on page 115

Unexpected Denial-of-Service Protection Behavior During Unified ISSU

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.

Related Documentation

- Application Support for Unified ISSU on page 115

Unexpected Ethernet Behavior During Unified ISSU

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 JunosE 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, 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.

Related Documentation

- Application Support for Unified ISSU on page 115

Unexpected File Transfer Protocol Server Behavior During Unified ISSU

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 progress	
	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.

Related Documentation

- Application Support for Unified ISSU on page 115

IS-IS Effects on Graceful Restart and Network Stability During Unified ISSU

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 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.

Related Documentation

- Application Support for Unified ISSU on page 115

Unexpected L2TP Failover of Established Tunnels During Unified ISSU

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.

Related Documentation

- Application Support for Unified ISSU on page 115

OSPF Effects on Graceful Restart and Network Stability During Unified ISSU

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 Protocols During Unified ISSU” on page 135

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.

Related Documentation

- Application Support for Unified ISSU on page 115
- `overload advertise-high-metric issu`
- `issu start`

Unexpected Suspension of PIM 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.

Related Documentation

- Application Support for Unified ISSU on page 115
- `ip pim dr-priority`

Unexpected Suspension of Subscriber Login and Logouts 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.

Related Documentation

- Application Support for Unified ISSU on page 115

Unexpected SONET and 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,

**Related
Documentation**

- Application Support for Unified ISSU on page 115

Unexpected T3 Behavior During Unified ISSU

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 JunosE 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.

**Related
Documentation**

- Application Support for Unified ISSU on page 115

Unavailability of TACACS+ Services During Unified ISSU

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.

**Related
Documentation**

- Application Support for Unified ISSU on page 115

Interruption in Traffic Forwarding for Layer 3 Routing Protocols During Unified ISSU

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 21 on page 136 describes how individual routing protocols behave during a unified in-service software upgrade.

Table 21: 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.

Table 21: Behavior of Routing Protocols During a Unified In-Service Software Upgrade *(continued)*

Protocol	Behavior
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.

Related Documentation

- Application Support for Unified ISSU on page 115

Recommended Settings for Routing Protocol Timers During Unified ISSU

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 22 on page 138.

Table 22: 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 Protocols During Unified ISSU" on page 135.</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 Protocols During Unified ISSU" on page 135.</p>

Table 22: 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.

Related Documentation • Application Support for Unified ISSU on page 115

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 *JunosE 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 JunosE 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
```

Related Documentation

- [issu initialize](#)
- [issu start](#)
- [issu stop](#)
- [license issu](#)

Halt of Unified ISSU During Initialization Phase Overview

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.

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

Related Documentation

- **issu stop**
- Halting Unified ISSU During Initialization Phase on page 142
- Halt of Unified ISSU During Upgrade Phase Overview on page 143
- Halting Unified ISSU During Upgrade Phase on page 143

Halting Unified ISSU During Initialization Phase

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
```

Related Documentation

- Halt of Unified ISSU During Initialization Phase Overview on page 142

- `disable-autosync`
- `boot system`
- `synchronize`
- `show redundancy`

Halt of Unified ISSU During Upgrade Phase Overview

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

Halting Unified ISSU During Upgrade Phase

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
```

- 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.

Related Documentation

- Unified ISSU Phases Overview on page 107
- Halt of Unified ISSU During Upgrade Phase Overview on page 143
- disable-autosync
- boot system
- synchronize

Monitoring the Status of the Router During Unified ISSU

Purpose Display information about the current status of the router relative to a unified in-service software upgrade and of the upgrade itself.

Action To display the current unified ISSU state and identify the active and armed releases in brief:

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
```

To the information displayed by **show issu brief**, the **show issu** command 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

To the information displayed by **show issu**, the **show issu detail** command 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

Meaning Table 23 on page 146 lists the **show issu** command output fields.

Table 23: show issu Output Fields

Field Name	Field Description
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 JunosE Software release that is currently running on the SRP modules
armed release	Filename of JunosE Software release that is armed to become the next running release when the router reboots
ISSU Activation Criteria Summary	Summarizes the criteria for unified ISSU activation.
In-Service Software Upgrade ready?	Criteria required for unified ISSU activation. Possible values: Yes, No, Conditional. NOTE: All criteria must be "yes" for unified ISSU to be active.
High-Availability ready?	Criteria required for unified ISSU activation. Possible values: Yes, No, Conditional. NOTE: All criteria must be "yes" for unified ISSU to be active.
Line Modules ready?	Criteria required for unified ISSU activation. Possible values: Yes, No, Conditional. NOTE: All criteria must be "yes" for unified ISSU to be active.
Configuration conversion support ready?	Criteria required for unified ISSU activation. Possible values: Yes, No, Conditional. NOTE: All criteria must be "yes" for unified ISSU to be active.

Table 23: show issu Output Fields (*continued*)

Field Name	Field Description
CLI sessions ready?	Criteria required for unified ISSU activation. Possible values: Yes, No, Conditional. NOTE: All criteria must be "yes" for unified ISSU to be active.
Routing applications ready?	Criteria required for unified ISSU activation. Possible values: Yes, No, Conditional. NOTE: All criteria must be "yes" for unified ISSU to be active.
Protocol Timers ready?	Criteria required for unified ISSU activation. Possible values: Yes, No, Conditional. NOTE: All criteria must be "yes" for unified ISSU to be active.
ISSU Criterion Detail	Detailed information on why the criteria required for unified ISSU activation was not met or was conditional.
Problem	The reason why the criteria for unified ISSU activation is not met.
Reporting Slot	Slot where the issue occurred.
Impact	What happens if you continue with the upgrade.
Remedy	What you can do to fix the Problem.

Related Documentation

- [show issu](#)

CHAPTER 6

Configuring VRRP

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

- VRRP Overview on page 149
- VRRP Platform Considerations on page 150
- VRRP Terms on page 151
- VRRP References on page 151
- VRRP Implementation in E Series Routers on page 152
- VRRP Router Election Rules on page 152
- Example: Basic VRRP Configuration on page 153
- Example: Commonly Used VRRP Configuration on page 154
- Example: VRRP Configuration Without the Real Address Owner on page 155
- Before You Configure VRRP on page 156
- Configuring VRRP on page 157
- Changing the Object Priority on page 159
- Monitoring the Configuration of VRIDs on page 159
- Monitoring the Configuration of VRRP Neighbors on page 162
- Monitoring the Statistics of VRRP Routers on page 163
- Monitoring the Configuration of VRRP Tracked Objects on page 166

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 *JunosE 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.

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.

**Related
Documentation**

- VRRP Implementation in E Series Routers on page 152
- VRRP Router Election Rules on page 152
- Before You Configure VRRP on page 156
- Configuring VRRP on page 157

VRRP 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.

VRRP Terms

Table 24 on page 151 provides definitions for the basic VRRP terms used in this chapter.

Table 24: 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.

VRRP References

For more information about VRRP, see:

- 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 *JunosE IP, IPv6, and IGP Configuration Guide*.

- RFC 3768—Virtual Router Redundancy Protocol (VRRP) (April 2004)

VRRP Implementation 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 3768 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.

Related Documentation

- VRRP Router Election Rules on page 152
- Before You Configure VRRP on page 156
- Configuring VRRP on page 157
- `ip vrrp accept-data`

VRRP 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.

Related Documentation

- VRRP Implementation in E Series Routers on page 152
- Before You Configure VRRP on page 156
- Configuring VRRP on page 157


Example: Basic VRRP Configuration

As Figure 6 on page 154 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 6: Basic VRRP Configuration




Related Documentation

- VRRP Overview on page 149
- VRRP Implementation in E Series Routers on page 152
- VRRP Router Election Rules on page 152
- Before You Configure VRRP on page 156
- Configuring VRRP on page 157

Example: Commonly Used VRRP Configuration

Figure 7 on page 155 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 7: Commonly Used VRRP Configuration


**Related Documentation**

- VRRP Overview on page 149
- VRRP Implementation in E Series Routers on page 152
- VRRP Router Election Rules on page 152
- Before You Configure VRRP on page 156
- Configuring VRRP on page 157

Example: VRRP Configuration Without the Real Address Owner

Figure 8 on page 156 is noticeably similar to Figure 7 on page 155 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 8: 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.

Related Documentation

- VRRP Overview on page 149
- VRRP Implementation in E Series Routers on page 152
- VRRP Router Election Rules on page 152
- Before You Configure VRRP on page 156
- Configuring VRRP on page 157

Before You Configure VRRP

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.

Related Documentation

- Configuring VRRP on page 157
- VRRP Overview on page 149

Configuring VRRP

Before you configure VRRP, we recommend that you review the following VRRP configuration examples:

- Example: Basic VRRP Configuration on page 153
- Example: Commonly Used VRRP Configuration on page 154
- Example: VRRP Configuration Without the Real Address Owner on page 155

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) Enable the backup router to learn the VRRP advertisement interval.

```
host1(config-if)#ip vrrp 22 timers-learn
```

5. (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
```

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

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

7. Associate an IP address with a VRID.

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



NOTE: If you configure VRRP on a virtual router and associate the IP address with the VRRP instance ID (VRID) so that the virtual address becomes the interface address of the router, the priority of the router automatically changes to 255 making it the master router. This change of priority occurs in JunosE Software Releases 11.0.0 and higher-numbered releases and later to enable full compliance with RFC-Virtual Router Redundancy Protocol (VRRP) (April 2004).

Also, you cannot configure the priority of the VRRP router as 255 by using the `ip vrrp priority` command, unless you configured the router to automatically learn associated addresses by using the `auto` keyword with the `ip vrrp virtual-address` command. In addition, if you change the virtual address of the VRRP router, which is operating as the IP address owner, to an IP address that is no longer the IP address owner, the priority changes automatically to the default value of 100.

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

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

9. (Optional) Configure the VRRP authentication key.

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

10. Enable the VRID instance.

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

Related Documentation

- Before You Configure VRRP on page 156
- VRRP Overview on page 149
- VRRP Implementation in E Series Routers on page 152
- `ip vrrp`
- `ip vrrp accept-data`
- `ip vrrp advertise-interval`
- `ip vrrp authentication-key`
- `ip vrrp authentication-type`
- `ip vrrp enable`
- `ip vrrp preempt`
- `ip vrrp priority`

- ip vrrp timers-learn
- ip vrrp virtual-address

Changing the 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.

To dynamically change the priority of a virtual router ID (VRID) in response to a change in the state of a specified object:

1. Track an object by its virtual ID. This example creates a new VRID.

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



NOTE: Multiple VRIDs can track the same object and a single VRID can track multiple objects.

2. Specify the value by which the priority must be decremented. This example dynamically changes the priority of the VRID in response to a change in state of object abc.

```
host1(config-if)#ip vrrp 25 track abc decrement 15
```



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

Related Documentation

- ip vrrp track

Monitoring the Configuration of VRIDs

Purpose Display information about all configured VRIDs.

Action To display a detailed summary of all configured VRIDs:

```
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

```

To display the summary count on all configured VRIDs:

```

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

```

To display a brief summary of all configured VRIDs:

```

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

```

Meaning Table 25 on page 160 lists the **show ip vrrp** command and **show ip vrrp summary** command output fields.

Table 25: show ip vrrp and show ip vrrp summary Output Fields

Field Name	Field Description
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: <ul style="list-style-type: none"> master—Router that forwards packets sent to the IP address associated with the virtual router. backup—Router that forwards packets if the current master router fails. Provides the current master router's IP address. init—Router that transitions to either the master state or backup state depending on the priority assigned.
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

Table 25: show ip vrrp and show ip vrrp summary Output Fields (*continued*)

Field Name	Field Description
Learning timer mode	Mode of the VRRP router: <ul style="list-style-type: none"> enabled—Router learns the VRRP advertisement interval that is useful in case of failure of the master router disabled—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	Type of authentication used by VRRP: none or text
preemption	Status of VRRP router preemption: enabled or disabled <ul style="list-style-type: none"> enabled—Backup router always takes over the responsibility of the master router. disabled—Backup router with lower-priority remains in backup state.
accept data	Accept data status of the VRRP router: <ul style="list-style-type: none"> enabled—Enables 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. disabled—Disables the processing of data packets by the backup router while the router is in the master state.
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

Table 25: show ip vrrp and show ip vrrp summary Output Fields (continued)

Field Name	Field Description
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
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 the VRRP router: <ul style="list-style-type: none"> • master—Router that forwards packets sent to the IP address associated with the virtual router. • backup—Router that forwards packets if the current master router fails. Provides the current master router's IP address. • init—Router that transitions to either the master or backup router depending on the priority assigned.
Adv	Advertisement interval, in seconds
Pri	Priority assigned to this router
Admin	Administrative state of the VRID: enabled or disabled

- Related Documentation**
- VRRP Overview on page 149
 - Before You Configure VRRP on page 156
 - Configuring VRRP on page 157

Monitoring the Configuration of VRRP Neighbors

Purpose Display neighbor information to the VRRP routers. Neighbor is a router that shares a given VRID with the VRRP router. A neighbor is known to the VRRP router only when the neighbor becomes a master for an IP address and sends VRRP advertisements. If a router sharing the VRID has not yet become a master router, then the local router remains unaware of this neighbor and this command does not display that neighbor.

Action To display information about all known neighbors to the VRRP routers:

```
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

```

Meaning Table 26 on page 163 lists the **show ip vrrp neighbor** command output fields.

Table 26: show ip vrrp neighbor Output Fields

Field Name	Field Description
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. If the priority value is 255, then the VRRP router is the master.
auth type	VRRP authentication type: none or text
assoc address(es)	IP addresses associated with the VRID that are advertised by the neighbor

- Related Documentation**
- VRRP Overview on page 149
 - Before You Configure VRRP on page 156
 - Configuring VRRP on page 157

Monitoring the Statistics of VRRP Routers

Purpose Display global statistics, interface statistics, or statistics per interface and VRID of configured VRRP routers.

Action To display the 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
```

To display the 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
```

To display the global statistics of a VRRP router:

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

Meaning Table 27 on page 165 lists the **show ip vrrp statistics** command output fields.

Table 27: show ip vrrp statistics Output Fields

Field Name	Field Description
Globals	
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

Table 27: show ip vrrp statistics Output Fields (*continued*)

Field Name	Field Description
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

- Related Documentation**
- VRRP Overview on page 149
 - Before You Configure VRRP on page 156
 - Configuring VRRP on page 157

Monitoring the Configuration of VRRP Tracked Objects

Purpose Display details of objects tracked by various VRIDs.

Action To display the details of objects tracked using VRIDs.

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

Meaning Table 28 on page 166 lists the **show ip vrrp tracked-objects** command output fields.

Table 28: show ip vrrp tracked-objects Output Fields

Field Name	Field Description
Interface	Name of the Interface. Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface specifier.
Vrid	VRRP router instance configured on the interface
Priority	Priority of the VRRP router. If the priority value is 255, then the VRRP router is the master.
Object	Name of the object being tracked
Type	Type of object being tracked

Table 28: show ip vrrp tracked-objects Output Fields (*continued*)

Field Name	Field Description
State	State of the object
Decrement	Value by which the priority is decremented or restored following an object state change

**Related
Documentation**

- Changing the Object Priority on page 159

CHAPTER 7

Managing Interchassis Redundancy

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

- ICR Overview on page 169
- ICR Platform Considerations on page 171
- ICR Terms on page 172
- ICR References on page 172
- ICR Scaling Considerations on page 172
- Interaction with RADIUS for ICR on page 173
- Configuring an ICR Partition on page 175
- Configuring the Interface on Which the ICR Partition Resides on page 176
- Configuring VRRP Instances to Match ICR Requirements on page 177
- Naming ICR Partitions on page 178
- Grouping ICR Subscribers Based on S-VLAN IDs on page 178
- Grouping ICR Subscribers Based on VLAN IDs on page 179
- Example: Configuring ICR Partitions That Group Subscribers by S-VLAN ID on page 180
- Using RADIUS to Manage Subscribers Logging In to ICR Partitions on page 182
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 183
- Monitoring the Configuration of ICR Partitions on page 184

ICR Overview


A broadband services router (BSR) aggregates many subscribers and services such as video on demand (VoD), voice over IP (VoIP), Internet Protocol television (IPTV), and the Internet, simultaneously. If the router fails because of hardware failures, subscriber downtime can result.

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

Virtual Router Redundancy Protocol (VRRP) to detect failures. ICR also enables you to track the failure of uplink interfaces. ICR currently supports only PPPoE subscribers.

Figure 9 on page 170 illustrates ICR deployment.

Figure 9: ICR Deployment



The subscriber broadcasts a PPPoE Active Discovery Initiation (PADI) packet to both the *master* and *backup* router. Only the *master* router processes the packet and creates the subscriber session. When the *master* router fails, VRRP switchover occurs and the *backup* router becomes the new *master* router. When receiving traffic for non-existent PPPoE sessions, the new *master* router sends early termination requests by sending PPPoE Active Discovery Termination (PADT) packets to the clients instead of waiting for the client to reconnect after the PPPoE session expires. The clients respond by sending requests to log in again. Then, the new *master* router creates new sessions for the PPPoE subscribers.

In lower-numbered releases, the new *master* router dropped the PPPoE packets because a session did not exist for the PPPoE subscribers and did not send PADT packets.

ICR achieves load balancing in case of failures on a per physical port basis by enabling you to create partitions. An *ICR partition* is a set of S-VLANs (and CVLANs) associated with a unique VRRP instance. There can be multiple partitions per physical port. A partition is the basic unit of redundancy. A partition cannot span multiple physical ports.

You can also create ICR clusters. An *ICR cluster* consists of a group of routers 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 ensure that the aggregation layer between the E Series router and access node (DSLAM) provides a broadcast domain per VLAN or per S-VLAN between active and backup routers. In the case of a direct connect model the access node must provide the broadcast domain per VLAN or per S-VLAN between the active and backup routers or instead provide an Ethernet switch such as EX Series Ethernet Switch between the access node and E Series router.

- Related Documentation**
- ICR Scaling Considerations on page 172
 - Configuring an ICR Partition on page 175

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/adaptor/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/adaptor/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
Documentation**

- Interface Types and Specifiers

ICR Terms

Table 29 on page 172 defines terms used in this discussion of ICR.

Table 29: 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 or backup.
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

When planning an ICR cluster you must ensure that you have provisioned adequate backup capacity in the event of a worst-case failure scenario such as a multiple hardware or multiple router failure.



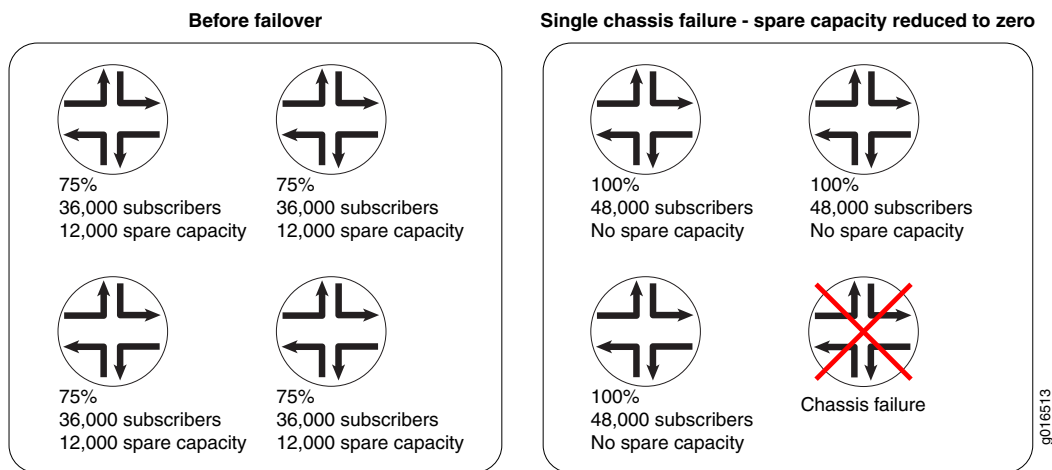
NOTE: Remember to consider parameters such as link bandwidth, QoS, and line module scaling limitations when you plan the deployment of the ICR cluster.

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 10 on page 173. 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 can be loaded with around 75 percent of its capacity while the remaining 25 percent is available 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 10 on page 173 illustrates an example of a typical ICR configuration.

Figure 10: Sample 1:1 Subscriber Redundancy in a 4–Node ICR Cluster



Related Documentation • [Configuring an ICR Partition on page 175](#)

Interaction with RADIUS for ICR

Authorization and authentication access messages identify subscribers before the RADIUS server grants or denies those subscribers access to the network or network services. When an application requests user authentication, the request must have certain authenticating attributes, such as a user's name, password, and the particular type of service the user is requesting. This information is sent in the authentication request via the RADIUS protocol to the RADIUS server. In response, the RADIUS server grants or denies the request.

JunosE Software supports certain RADIUS vendor-specific attributes (VSAs) that define specific authentication, authorization, and accounting elements in a user's profile. The profile is stored on the RADIUS server. RADIUS messages contain RADIUS attributes to communicate information between an E Series Broadband Services Router and the RADIUS server. For complete information on VSAs, see *Configuring RADIUS Attributes* in the *JunosE Broadband Access Configuration Guide*. JunosE Software Release 10.3.x and later supports the ICR-Partition-Id VSA [26-150]. You can use this VSA to collect information on the ICR partition configured on the VLAN or S-VLAN subinterface on which subscribers are logged in.

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 *JunosE Broadband Access Configuration Guide*.

Determining the ICR partition is useful for accounting and authentication of subscribers in RADIUS messages.

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.

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 corresponding ICR interface is configured.

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.

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 cannot 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 be 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 fails, the Partition-Accounting-Off message is sent from the failing B-RAS application on the router after Acct-Stop responses are received for all the logged out subscribers in that partition.

Related Documentation

- Using RADIUS to Manage Subscribers Logging In to ICR Partitions on page 182
- `radius include`
- `radius icr-partition-accounting`
- `show radius icr-partition-accounting`
- *Configuring RADIUS Attributes* in the *JunosE Broadband Access Configuration Guide*

Configuring an ICR Partition

You can use RADIUS servers to authenticate subscribers and collect statistics related to the users logging in to an ICR partition on a virtual router. 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 176.
2. Create a unique VRRP instance to manage the ICR partition.
See “Configuring VRRP Instances to Match ICR Requirements” on page 177.
3. Create and assign a name to the ICR partition.
See “Naming ICR Partitions” on page 178.
4. (Optional) Select the grouping criterion for the ICR partition.
See “Grouping ICR Subscribers Based on S-VLAN IDs” on page 178 and “Grouping ICR Subscribers Based on VLAN IDs” on page 179.



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 182.

**Related
Documentation**

- ICR Overview on page 169
- Monitoring the Configuration of ICR Partitions on page 184
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 183
- 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 *JunosE Release Notes, Appendix A, System Maximums*.

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
Documentation**

- Configuring VRRP Instances to Match ICR Requirements on page 177
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 183

Configuring VRRP Instances to Match ICR Requirements

Each ICR partition is managed by a unique VRRP instance. You can specify an ICR partition as the *master* partition by assigning a higher priority. Use the **ip vrrp priority** command to assign priorities to the ICR partitions.

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. (Optional) Enable the router to learn the VRRP advertisement interval. Use this only when you plan on upgrading your router by means of a unified in-service software upgrade (ISSU).

```
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.

**Related
Documentation**

- VRRP Overview on page 149
- 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.

```
host1(config-if)#ip vrrp 2 icr-partition part1Backup
```

```
host1(config-if)#ip vrrp 3 icr-partition ICRBackup
```

For information on the number of ICR partitions that you can create per line module or chassis, see *JunosE Release Notes, Appendix A, System Maximums*.

Related Documentation

- ICR Overview on page 169
- Monitoring the Configuration of ICR Partitions on page 184
- `ip vrrp icr-partition`

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 the state of the corresponding subinterfaces (up/AdminDown) based on the state of the partition (master or backup). If the subinterfaces are part of the backup partition, the router changes the state of all the subinterfaces to AdminDown.
- Use the **use-default-mac** keyword to enable the subinterfaces to use the default MAC address instead of the VRRP MAC address. By default, subinterfaces use the virtual MAC address of the associated VRRP instance.
- 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 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
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
advertise-mac
```



NOTE: To enable the new master router to send PPPoE Active Discovery Termination (PADT) packets to the clients and create new sessions for the PPPoE subscribers, you must create a dummy IP interface for each S-VLAN that is part of the ICR partition.

4. (Optional) Configure additional S-VLAN subinterfaces by completing Steps 2 and 3 using unique numbering.

Related Documentation

- Grouping ICR Subscribers Based on VLAN IDs on page 179
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 183
- ip vrrp icr-partition group
- ip vrrp icr-partition svlan-list
- ip vrrp icr-partition svlan-list explicit
- ip vrrp icr-partition svlan-range

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 the state of the corresponding subinterfaces (up/AdminDown) based on the state of the partition (master or backup). If the subinterfaces are part of the backup partition, the router changes the state of all the subinterfaces to AdminDown.

- Use the **use-default-mac** keyword to enable the subinterfaces to use the default MAC address instead of the VRRP MAC address. By default, subinterfaces use the virtual MAC address of the associated VRRP instance.
- 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  
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  
advertise-mac
```

3. (Optional) Configure additional VLAN subinterfaces by completing Step 2 using unique numbering.

Related Documentation

- Grouping ICR Subscribers Based on S-VLAN IDs on page 178
- Monitoring the Configuration of an ICR Partition Attached to an Interface on page 183
- `ip vrrp icr-partition group`
- `ip vrrp icr-partition vlan-list`
- `ip vrrp icr-partition vlan-range`

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.

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-interface
host1 (config-if)#ip vrrp 1 icr-partition svlan-range 111 119 advertise-mac
host1 (config-if)#ip vrrp 1 icr-partition svlan-list-explicit 120 1 120 2 advertise-mac
control-interface
host1 (config-if)#exit

```



NOTE: To enable the new master router to send PPPoE Active Discovery Termination (PADT) packets to the clients and create new sessions for the PPPoE subscribers, you must create a dummy IP interface for each S-VLAN that is part of the ICR partition.

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-interface
host2 (config-if)#ip vrrp 1 icr-partition svlan-range 111 119 advertise-mac
host2 (config-if)#ip vrrp 1 icr-partition svlan-list-explicit 120 1 120 2 advertise-mac
control-interface
host2 (config-if)#exit
```



NOTE: To enable the new master router to send PPPoE Active Discovery Termination (PADT) packets to the clients and create new sessions for the PPPoE subscribers, you must create a dummy IP interface for each S-VLAN that is part of the ICR partition.

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
Documentation**

- ICR Overview on page 169
- ICR Scaling Considerations on page 172

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 *JunosE 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 *JunosE Broadband Access Configuration Guide*.

- Related Documentation**
- Interaction with RADIUS for ICR on page 173
 - Configuring an ICR Partition on page 175
 - radius include
 - radius icr-partition-accounting
 - show radius icr-partition-accounting

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 3/5/0.11`
 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 30 on page 183 lists the `show icr-partition` command output fields.

Table 30: 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>

Table 30: show icr-partition Output Fields (*continued*)

Field Name	Field Description
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.
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 the state of the corresponding subinterfaces (up/AdminDown) based on the state of the partition (master or backup). If the subinterfaces are part of the backup partition, the router changes the state of all the subinterfaces to AdminDown. You can also block all traffic on the backup partition. However, the router does not block VRRP advertisements as long as VRRP is running on a separate interface. Possible states: enabled or disabled. If the status is enabled, the router changes the state of the subinterface based on the state of the partition. If the status is disabled, the router does not control the state of the corresponding subinterface.
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 VRRP MAC address; otherwise, the interface uses the default 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 Documentation**
- Configuring the Interface on Which the ICR Partition Resides on page 176
 - 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 a specific 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 31 on page 185 lists the **show icr-partitions** command output fields.

Table 31: 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.

Table 31: show icr-partitions Output Fields (*continued*)

Field Name	Field Description
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>
Partition-ID	Identifier for the ICR partition.
Dormant ICR Partitions	Number of dormant ICR partitions 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
Documentation**

- [Configuring the Interface on Which the ICR Partition Resides](#) on page 176
- [show icr-partitions](#)

PART 2

Index

- Index on page 189

Index

A

access modules	
ES2 4G LMs as	
in an LNS device, stateful switchover.....	72
access modules in the LNS	
receipt of event from an application	
access module processes PPP echo	
requests.....	81
failure with the primary module.....	81
Access-Request messages	
ICR Partition ID VSA.....	174
Acct-Start messages	
ICR Partition ID VSA.....	174
Acct-Stop messages	
ICR Partition ID VSA.....	174
activating guidelines	
high availability, line modules	
mirroring configuration files from active to	
standby.....	87
reverting to default redundancy mode.....	87
same software release on primary and	
secondary modules.....	87
active	
high availability state.....	40
active, line module high availability state	
change in previously matching criterion	
transition to disabled state.....	87
mirroring of information	
relevant applications create and post	
transactions.....	86
standby module replays updates and	
performs changes for applications.....	86
updates in settings sent to standby	
module.....	86
synchronization of data	
from active to standby module.....	86
unsupported application configured	
transition to disabled state.....	87
when a switchover occurs	
warm restart of standby module.....	87

application support	
high availability.....	42
Automatic Reversion	
enabling.....	13
automatic switchover.....	12
limitations.....	12
Automatic Switchover	
disabling.....	13

B

backup router.....	152
defined.....	151
election process and.....	152
VRRP.....	149

C

call setup rate	
data synchronization	
impact on.....	72
under peak load conditions	
time for system to become HA-active.....	72
clear redundancy commands	
clear redundancy history.....	65
cold boot	
of secondary line module	
when slot is disabled.....	78
when slot is reloaded.....	78
when software fault occurs.....	78
commands for configuring	
stateful SRP switchover	
similarity with commands for stateful line	
module switchover.....	69
connection manager	
usage of stream ID	
to identify connections.....	80
conventions	
notice icons.....	xix
text and syntax.....	xx
customer support.....	xxi
contacting JTAC.....	xxi

D

data synchronization	
call setup rate impact	
consumption of backplane bandwidth.....	72
deactivating	
high availability, line modules.....	89
deactivating guidelines	
high availability, line modules.....	89
redundancy mode of line modules	
disruption in user sessions, forwarding data.....	89
reloading interfaces, routing tables from SRP.....	89
destination address (DA), VRRP.....	153
direct memory access (DMA)	
receipt of packets for forwarding controller	
by interface controller.....	75
receipt of packets for system controller	
by interface controller.....	75
disable-switch-on-error command.....	21
disabled	
high availability state.....	39
disabled, line module high availability state	
cold-restart of the router	
when switchover occurs.....	85
retention in the same state	
requirements not met.....	85
disabled,, line module high availability state	
criteria to move to initializing state	
active module support for high availability.....	85
high availability mode configured.....	85
releases on standby and active modules are same.....	85
standby module is online, mirroring enabled.....	85
standby module support for high availability.....	85
router uses redundancy mode.....	85
documentation set	
comments on.....	xxi
downlink modules	
ES2 4G LMs as	
in an LNS device, stateful switchover.....	72
drop events	
on egress queues	
not preserving across stateful switchover of LMs.....	72

drop rates	
on egress queues	
not preserving across stateful switchover of LMs.....	72

E

egress queues	
drop events on	
stateful switchover of LMs.....	72
drop rates on	
stateful switchover of LMs.....	72
forwarding events on	
stateful switchover of LMs.....	72
forwarding rates on	
stateful switchover of LMs.....	72
ES2 4G line modules	
as access modules	
in an LNS device for stateful switchover.....	72
as downlink modules	
in an LNS device for stateful switchover.....	72
ES2 4G LMs	
installed with Service IOA	
support for stateful line module switchover.....	68

F

failover. See switchover	
file system synchronization mode	
redundancy mode.....	37
forwarding controller	
receives packets from interface controller	
using direct memory access.....	75
tables that point to failed modules	
updated with stream IDs to new primary.....	81
transfer of packets from the interface controller	
not preserved during stateful line module switchover.....	71
transfer of packets to the system controller	
not preserved during stateful line module switchover.....	71
usage of stream ID	
to identify connections from a specific slot.....	80
forwarding controller database	
mapping of the slot ID, stream ID, and traffic class.....	80

forwarding events	
on egress queues	
not preserving across stateful switchover	
of LMs.....	72
forwarding rates	
on egress queues	
not preserving across stateful switchover	
of LMs.....	72

H

hardware	
monitoring information.....	23
high availability	
activating.....	20, 52
activating guidelines.....	51
deactivating.....	53
IP interface priority.....	54
overview.....	35
high availability mode	
redundancy mode.....	37
high availability mode, line modules	
brief pause in data forwarding	
when secondary module becomes	
primary.....	83
initial bulk transfer of data	
to newly assigned primary module.....	83
no impact on existing user sessions.....	83
synchronization of state and dynamic	
configuration	
between primary and secondary LMs.....	83
transaction-based mirroring subsequently	
to newly assigned primary module.....	83
high availability pair	
of line modules	
unified ISSU disabled on the	
secondary.....	71
unified ISSU support on the primary.....	71
high availability, line module See stateful line	
module switchover	
high availability, line modules	
activating	
launching Redundancy Configuration	
mode.....	88
specifying slots of primary and secondary	
modules.....	88
activating guidelines.....	87
deactivating.....	89
deactivating guidelines.....	89

error message displayed	
when maximum limit is exceeded for pairs	
configured.....	88
maximum pairs that can be configured.....	88

I

icr cluster.....	172
ICR commands	
interface	176
icr interface.....	172
ICR Options	
icr-control-interface.....	177
priority command.....	177
timers-learn command.....	177
ICR Partition	
configuring	175, 176, 177, 178, 179
configuring, naming.....	178
radius.....	182
ICR partition accounting	
and dependence on Acct-Stop messages.....	175
configuring.....	173
disabling and enabling messages	
sent to the RADIUS server.....	173
overview.....	173
processing in different scenarios	
administrative switchover.....	175
complete chassis failure.....	175
line module or interface failure.....	175
transition of ICR partition states.....	175
ICR Partition commands	
naming	178
ICR partition ID VSA	
including in access and accounting	
messages.....	174
ICR Partition ID VSA	
transmitting to the virtual router	
where ICR control interface is	
configured.....	174
where ICR partition is configured.....	174
ICR Partition Options	
advertise-mac.....	179, 180
control-interface	179, 180
group option.....	178, 180
use-default-mac	179, 180

ICR Partitions		
configuration example		
backup ICR partition, S-VLAN based		
grouping.....	181	
master ICR partition, S-VLAN based		
grouping.....	180	
ICR RADIUS commands		
inclusion of icr-partition-id.....	182	
radius icr-partition-accounting.....	182	
icr-partition.....	172	
in-service software upgrade. <i>See</i> unified ISSU		
initializing		
high availability state.....	40	
initializing, line module high availability state		
completion of operation		
transition to active state.....	86	
criteria not met or previously matching criteria		
not complied		
shifts to disabled state.....	86	
events that occur		
bulk synchronization of memory.....	86	
mirroring of state and dynamic		
configuration.....	86	
presence of unsupported application		
shifts to disabled state.....	86	
when a switchover occurs		
system cold-restarts.....	86	
Interchassis Redundancy		
heterogenous icr clusters.....	170	
icr clusters.....	170	
icr partition.....	170	
interface controller		
forwards packets to system controller		
using direct memory access.....	75	
using Ethernet channel.....	75	
operational image that runs on		
and stateful line module switchover.....	70	
receives packets destined for forwarding		
controller		
using Ethernet channel.....	75	
receives packets destined for system controller		
using direct memory access.....	75	
Interim-Acct messages		
ICR Partition ID VSA.....	174	
IOA slots		
and SRP module combination		
compatible with stateful line module		
switchover.....	70	
IP addresses		
IP address owner, VRRP.....	151	
primary, VRRP.....	151	
VRRP.....	152, 157	
ip commands.....	157	
ip address.....	156	
ip vrrp.....	157	
ip vrrp accept-data.....	157	
ip vrrp advertise-interval.....	157	
ip vrrp authentication-key.....	157	
ip vrrp authentication-type.....	157	
ip vrrp enable.....	157	
ip vrrp preempt.....	157	
ip vrrp priority.....	157	
ip vrrp track.....	159	
ip vrrp virtual-address.....	157	
<i>See also</i> vrrp commands		
ip pim commands		
ip pm dr-priority.....	133	
ISSU. <i>See</i> unified ISSU		
L		
L2TP		
after switchover to the secondary module		
occurs		
restores operation data to the new		
primary.....	81	
configuration and operation data		
maintained in the line module.....	81	
mirrored to the standby module.....	81	
on the SRP module		
handles line module events.....	81	
restoration of configuration to the new primary		
similarity with warm start during unified		
ISSU.....	81	
L2TP tunnels		
data packets received on		
preventing requests from LAC devices for		
sequence numbering.....	72	
sequence number checking for.....	72	
LEDs		
monitoring status.....	22	
line module high availability <i>See</i> stateful line		
module switchover		
line module high availability states		
active.....	84	
disabled.....	84	
initializing.....	84	

- line module redundancy
 - configuring.....13
 - managing.....13
 - modules enabled for
 - cannot be configured in a high availability pair.....71
 - overview.....9
 - requirements
 - E120 and E320 Routers.....10
 - ERX7xx models and ERX14xx Models.....10
- line modules
 - applications that support
 - stateful switchover.....79
 - behavior of system functions
 - with stateful switchover configured.....75
 - cold switchover of
 - loss of existign subscriber sessions.....78
 - same behavior as line module redundancy.....78
 - on E120 and E320 routers, components
 - forwarding controller.....74
 - input/output adapter.....74
 - interface controller.....74
 - replacing
 - when stateful switchover is enabled.....77
 - stateful switchover of
 - active or primary module.....68
 - packet processing by the router.....74
 - standby or secondary module.....68
 - warm switchover of
 - preservation of subscriber sessions.....78
 - with downlink interface to the LAC, reloading
 - not retaining user sessions during stateful switchover.....71
 - not retaining user sessions with line module redundancy.....71
- LNS devices
 - ES2 4G LMs as access modules
 - and stateful line module switchover.....72
 - ES2 4G LMs as downlink modules
 - and stateful line module switchover.....72
- LNS sessions
 - stateful switchover of
 - for routers that act as LNS devices.....70
 - supported module and IOA combinations.....70
- log messages
 - generated for stateful line module switchover
 - in response to manual settings.....68
 - in response to system events.....68
 - stateful line module switchover
 - transition between switchover states.....90, 91
 - transition from active to disabled state.....91
 - transition from disabled to active state.....91
 - when HA is enabled.....93
 - stateful SRP switchover
 - transition from pending to active state.....92
 - when HA is enabled.....92, 93
- M**
 - manuals
 - comments on.....xxi
 - master router.....151
 - Module Redundancy
 - switching example
 - primary Line Module, Spare Line Module.....14
 - modules
 - monitoring information.....27, 28
- N**
 - notice icons.....xix
- P**
 - pending
 - high availability state.....41
 - performance impact
 - maximum tunneled PPP sessions
 - on the primary line module.....72
 - platform considerations
 - high availability.....36
 - stateful line module switchover.....70
 - policy application
 - downloading policy attachments
 - using the SRP module.....79
 - in the line module
 - policy definitions not stored in.....79
 - transmission of policy attachments from SRP to LM
 - in a bulk manner.....79

policy attachments	
transfer in a bulk operation	
from SRP to line module.....	79
reduced time for download.....	79
policy manager	
stateful line module switchover safe	
download of policy attachments.....	72
preservation of statistics after	
switchover.....	72
policy statistics	
preserved across line module switchover	
green bytes and packets.....	94
red bytes and packets.....	94
saturated bytes and packets.....	94
transmit-unconditional bytes and	
packets.....	94
upper green bytes and packets.....	94
upper red bytes and packets.....	94
upper yellow bytes and packets.....	94
yellow bytes and packets.....	94
PPP accounting statistics	
preserved across line module switchover	
Acct-Input-Octets RADIUS attribute.....	93
Acct-Input-Packets RADIUS attribute.....	93
Acct-Output-Octets RADIUS	
attribute.....	93
Acct-Output-Packets RADIUS	
attribute.....	93
IPv6-Acct-Input-Octets RADIUS	
attribute.....	93
IPv6-Acct-Input-Packets RADIUS	
attribute.....	93
IPv6-Acct-Output-Octets RADIUS	
attribute.....	93
IPv6-Acct-Output-Packets RADIUS	
attribute.....	93
PPP application	
after stateful line module switchover	
replication of sessions on standby	
module.....	80
components on the line module	
basic protocol.....	80
state machines in a running state.....	80
timers.....	80
echo requests for sessions on ES2 4G LM	
handled by access module.....	80
echo requests for sessions on failed module	
redirection to a different hardware.....	80
expiry of keepalives	
time for the new primary to become	
active.....	80
mirrored storage data	
reconstruction of.....	80
sessions alternating between up and down	
states	
not retained during switchover.....	80
time for standby module to become active	
dependent on configuration settings.....	80
PPP echo reply messages	
sent from access modules in LNS	
in response to echo requests from	
clients.....	81
PPP echo requests	
handling by access module in LNS	
during stateful switchover.....	81
stoppage of handling by access module in LNS	
after stateful switchover is complete.....	81
PPP subscriber sessions	
on an LNS device in L2TP tunnels	
echo requests handled by access	
module.....	81
terminated due to lack of keepalive	
responses.....	81
primary line module	
actions that trigger stateful switchover	
disabling the slot.....	73
graceful switchover to secondary	
module.....	73
reloading the slot.....	73
collection of statistics from	
for policy manager.....	72
for PPP applications.....	72
disabling the slot.....	78
disabling the slot of	
secondary module becomes the	
primary.....	77
erasing the slot of	
primary and secondary modules are cold	
booted.....	78
failure of	
retaining active user sessions.....	69
running of diagnostic functions.....	69
stateful switchover to secondary	
module.....	69
subscriber disconnection for two	
minutes.....	69

- performance impact, stateful switchover
 - maximum tunneled PPP sessions.....72
 - reloading the slot of
 - secondary module becomes the
 - primary.....77
 - software fault occurs on
 - secondary module becomes the
 - primary.....77
 - timeout value
 - to reach the online state after router
 - reboot.....76
- Q**
- QoS**
- agent clients bind and register
 - to the QoS agent.....80
 - configuration stored in line modules.....80
 - settings mirrored to standby module.....80
 - stateful line module switchover support
 - restoration of queues on subscriber
 - interfaces.....72
- queue manager**
- agent running on line modules
 - handling of forwarding controller
 - updates.....80
 - agents reside on line modules.....80
 - initiation of requests to the connection
 - manager
 - on reception of the controller up
 - event.....80
 - resides on the SRP.....80
 - usage of queue ID
 - to identify connections.....80
- R**
- RADIUS**.....172
- rebooting of the router**
- stateful line module switchover configured
 - primary module reaches online state.....76
 - secondary module is not an online
 - state.....76
 - timeout for primary module to become
 - online.....76
- recovery of routers**
- from double failures
 - not supported.....71
- redundancy**
- line module. *See* line module redundancy
 - SRP module. *See* SRP module redundancy
- redundancy commands**
- redundancy force-switchover.....13, 21, 22
 - redundancy lockout.....13
 - redundancy revert.....13
 - redundancy revertive.....13
 - srp switch.....22
- redundancy modes**
- high availability.....37
- redundancy modes, line modules**
- high availability.....83
 - stateless switchover.....83
- removal of IOAs**
- without powering down
 - stateful line module switchover not
 - triggered.....78
- replacement of line modules**
- stateful line module switchover enabled, action
 - taken for failures
 - cold and warm switchovers of LMs in a HA
 - pair.....78
 - disabling the slots in which the LMs in a HA
 - pair reside.....78
 - reloading slot in which LMs in a HA pair
 - reside.....78
 - reloading the primary line module.....77
 - reloading the secondary line module.....77
 - removing IOAs from LMs in a HA pair
 - without powering down.....78
- resetting of subscriber sessions**
- with stateful line module switchover disabled,
 - dependency on
 - configuration parameters applied.....68
 - number of active user sessions.....68
 - router models in network.....68
- S**
- secondary line module**
- disabling the slot.....78
 - disabling the slot of
 - cold boot is performed.....78
 - erasing the slot of
 - line module high availability is
 - disabled.....78
 - mirroring of
 - subscriber details from failed primary.....69
 - performance impact, stateful switchover
 - newly provisioned.....72
 - preservation of active sessions
 - after a stateful switchover.....69

reloading the slot of	
cold boot is performed.....	78
software fault occurs on	
cold boot is performed.....	78
taking over as	
primary after stateful switchover.....	69
taking over as primary, timeout exceeded for	
primary	
after system reboot.....	76
taking over as the primary	
recovery of applications to a stable	
state.....	77
sequence number checking	
for packets on L2TP tunnels	
not maintained during stateful switchover	
of LMs.....	72
setting recommendation for stateful	
switchover of LMs.....	72
Service Availability	
Features.....	5
ICR.....	6
Module redundancy.....	5
Stateful Line Module Switchover.....	5
Stateful SRP Switchover.....	5
Unified ISSU.....	6
VRRP.....	6
show icr commands	
show icr-partition.....	183
show icr-partitions.....	184
state.....	185
summary.....	185
show ip commands	
show ip vrrp.....	159
show ip vrrp brief.....	159, 160
show ip vrrp neighbor.....	162
show ip vrrp neighbors.....	162
show ip vrrp statistics.....	163, 164
show ip vrrp statistics global.....	163, 164
show ip vrrp summary.....	159, 160
show ip vrrp tracked-objects.....	166
show issu commands	
show issu.....	144
show issu brief.....	144
show issu detail.....	144
show redundancy commands	
show environment.....	23
show hardware.....	27, 28
show redundancy.....	56
show redundancy clients.....	59
show redundancy clients all.....	59
show redundancy detail.....	56
show redundancy history.....	61
show redundancy history detail.....	61
show redundancy history line-card slot.....	97
show redundancy history line-card slot	
detail.....	97
show redundancy line-card.....	62
show redundancy line-card slot.....	95
show redundancy srp.....	63
show redundancy srp detail.....	63
show redundancy switchover-history.....	65
simultaneous switchover of	
SRP and line modules	
not supported.....	71
SRP module redundancy.....	16
managing.....	21
SRP modules	
installing a redundant module.....	19
paired with IOAs	
support for stateful line module	
switchover.....	70
reset button.....	16
simultaneous switchover with line modules	
not supported.....	71
stateful line module switchover	
1:1 redundancy model.....	70
access modules in an LNS device.....	72
actions that trigger	
disabling the slot of the primary	
module.....	73
graceful switchover to secondary	
module.....	73
rebooting the slot of the primary	
module.....	73
activating.....	88
activating guidelines.....	87
active module, primary.....	68
advantages of functionality	
1:1 redundancy model.....	68
seamless preservation of subscriber	
sessions.....	68
and line module redundancy	
cannot be enabled simultaneously.....	71
applications that do not support.....	79
applications that support	
connection manager.....	79
forwarding controller.....	79
interchassis control protocol.....	79

L2TP.....	79	forwarding controller tables	
mirroring subsystem.....	79	updated with stream IDs to new	
policy manager.....	79	primary.....	81
PPP.....	79	guidelines for configuring.....	70
QoS.....	79	historical record of, not preserving	
unified ISSU.....	79	drop events on egress queues.....	72
available on the operational image		drop rates on egress queues.....	72
that runs on the interface controller.....	70	forwarding events on egress queues.....	72
behavior of system functions		forwarding rates on egress queues.....	72
configured and disabled on LMs.....	75	hot-swap of the IOA	
configured and enabled on LMs.....	75	in a managed environment.....	78
interoperation with line module		improving router availability.....	68
redundancy.....	75	independent of stateful SRP switchover.....	71
interworking with stateful SRP		interworking with unified ISSU	
switchover.....	75	not supported on the secondary line	
LM is in active or disabled switchover		module.....	71
state.....	75	supported on the primary line module.....	71
when the router is rebooting.....	75	log messages	
when unified ISSU is performed.....	75	transition from active to disabled	
benefits of		state.....	91
enhanced reliability and resiliency.....	68	transition from disabled to active	
preservation of active client sessions.....	68	state.....	91
recovery of applications to a stable		log messages generated during	
state.....	68	in response to manual settings.....	68
undisrupted forwarding of data.....	68	in response to system events.....	68
deactivating.....	89	maximum pairs that can be configured.....	88
deactivating guidelines.....	89	modules enabled for	
differences in keywords of commands		cannot be configured in a redundancy	
with stateful SRP switchover.....	69	group.....	71
disabled		mutually exclusive of	
forcible termination of user sessions.....	69	line module redundancy.....	71
downlink modules in an LNS device.....	72	non-support of	
enabled on the router		1:N hot standby mode.....	71
retaining existing client connections.....	68	automatic switchover of serial	
synchronization of applications to a stable		connection.....	71
state.....	68	not preserving packets flowing	
uninterrupted forwarding on		between forwarding and interface	
interfaces.....	68	controllers.....	71
for LNS sessions		between forwarding and system	
active and standby Service IOAs		controllers.....	71
supported.....	70	overview.....	68
compatible SRP and SFM models.....	70	performance impact	
downlink and uplink LMs in an L2TP		insertion of new secondary module.....	72
tunnel.....	70	maximum PPP sessions on primary	
L2TP tunnels and sessions supported.....	70	module.....	72
with the router as an LNS in an L2TP			
tunnel.....	70		

primary line module	
continuation of data forwarding after	
failure of.....	69
failure of.....	69
mirroring of user sessions.....	69
scenarios of configuration.....	75
secondary line module	
mirroring of subscriber sessions from	
primary.....	69
not preserving fluctuating sessions.....	69
taking over as primary.....	69
sequence number checking	
for data packets on L2TP tunnels.....	72
similarity in configuration commands	
with stateful SRP switchover.....	69
spare or standby module, secondary.....	68
supported line module, IOA combination.....	69
supported LM and IOA combinations	
for stateful switchover of LNS	
sessions.....	70
system operations when enabled on the	
router.....	74
unavailable for setting up	
forcible termination of user sessions.....	69
using status LEDs for monitoring	
states of line modules.....	95
viewing	
detailed history information for line module	
in a slot.....	97
information about starts and switchovers	
for line module in a slot.....	97
redundancy setting of line module in a	
slot.....	95
stateful line module switchover commands	
line-card switch.....	73
line-card switchmode high-availability slot.....	73
stateful line module switchover, scenarios	
not supported in software releases	
restart of line modules.....	68
setting disabled on the router	
disconnection of user sessions.....	68
excessive time for rebuilding user	
sessions.....	68
reestablishment of client	
connections.....	68
restart of line modules.....	68
stateful SRP switchover	
configuration independent of	
stateful line module switchover.....	69
differences in keywords of commands	
with stateful line module switchover.....	69
disruption of connection between ICs	
reestablishment of ICC sessions.....	76
switch relearns MAC addresses.....	76
does not prevent root cause of reload or restart	
similarity in behavior with stateful line	
module switchover.....	70
log messages	
transition from pending to active	
state.....	92
similarity in configuration commands	
with stateful line module switchover.....	69
stateful SRP switchover.	37
<i>See also</i> high availability	
stateless switchover mode, line modules	
default behavior for redundant LMs.....	83
system events that occur.....	83
temporary halt in working of I/O modules.....	83
states	
high availability.....	38
statistics	
for policy manager	
collection frequency from primary	
module.....	72
polling and sending to secondary line	
module.....	72
for PPP applications	
collection frequency from primary	
module.....	72
polling and sending to secondary line	
module.....	72
last collected value used as baseline	
on the primary line module.....	72
status LEDs	
monitoring redundancy state of SRP and line	
modules	
OK LED.....	95
ONLINE LED.....	95
REDUNDANT LED.....	95
status LEDs, monitoring.....	22
subscriber sessions	
preservation of active	
during stateful line module	
switchover.....	68
support, technical <i>See</i> technical support	

switchover.....	10
stateless.....	9

T

technical support	
contacting JTAC.....	xxi
text and syntax conventions.....	xx
transmission of packets	
between forwarding and interface controllers	
not preserved during stateful line module	
switchover.....	71
between forwarding and system controllers	
not preserved during stateful line module	
switchover.....	71

U

unified ISSU	
disabled on the secondary module	
in a high availability pair.....	71
support on primary module	
in a high availability pair.....	71
unified ISSU (in-service software upgrade).....	101
AAA support.....	124
application support.....	115
ATM support.....	124
ATM port data rate.....	124
ILMI sessions.....	124
OAM CC effects.....	124
OAM VC integrity.....	124
VC and VP statistics.....	124
DHCP support.....	125
common component.....	125
external server.....	125
packet capture.....	125
relay and relay proxy.....	125
DHCP support:relay and relay proxy.....	116
DoS protection support.....	125
Ethernet support.....	126
ARP entries.....	126
LAG.....	126
port data rate.....	126
VLAN statistics.....	126
FTP support.....	127
halting during initialization.....	142
halting during upgrade.....	143

initialization phase.....	107
application data on standby SRP	
module.....	108
line module arming.....	108
SNMP traps.....	108
interoperation with	
stateful line module switchover.....	77
IS-IS support.....	129
graceful restart.....	129
high link cost.....	129
L2TP support.....	131
layer 3 protocol traffic forwarding.....	135
OSPF support.....	132
dead interval.....	132
graceful restart.....	132
high link cost.....	132
overview.....	102
phases	
initialization.....	108
overview of.....	107
service restoration.....	115
upgrade.....	110
PIM support.....	133
platform.....	104
procedure for upgrade.....	139
references.....	107
requirements	
hardware.....	104, 105
software.....	104
traffic forwarding.....	104
verification in upgrade phase.....	110
restoring original router state.....	142, 143
router behavior.....	102
service restoration phase.....	107, 115
SONET/SDH support.....	134
subscriber support.....	134
logins.....	134
statistics.....	134
support, application.....	115
T3 support.....	135
TACACS+ support.....	135
terms.....	106
timer settings for routing protocol timers.....	138
upgrade phase.....	107, 110
exceptions.....	110
line module control plane.....	110
line module forwarding plane	
upgrade.....	110
process steps.....	110

setup.....	110
SRP module switchover.....	110
verification requirements.....	110
upgrade procedure.....	139
unified ISSU (in-service software upgrade):DHCP support	
local server.....	116
upgrading software	
high availability.....	55

V

virtual MAC address.....	149
virtual router ID (VRID). See VRID	
Virtual Router Redundancy Protocol (VRRP). See VRRP	
VRID (virtual router ID)	
configuration.....	157
creating.....	157
router election rules.....	152
vrrp	
platform.....	150
VRRP.....	172
VRRP (Virtual Router Redundancy Protocol)	
advertisement interval.....	157
advertisement messages.....	152
authentication key.....	157
authentication type.....	157
backup router.....	149, 152
configuration examples.....	153, 154, 155
configuring.....	157
how it works.....	150
implementation.....	152
MAC address.....	149
master router.....	151
monitoring.....	159, 162, 163, 166
overview.....	149
preemption.....	152, 157
router election rules.....	152
router priority.....	157
VLAN support.....	149
VRRP router defined	151
vrrp commands	
ip vrrp.....	157
ip vrrp accept-data.....	157
ip vrrp advertise-interval.....	157
ip vrrp authentication-key.....	157
ip vrrp authentication-type.....	157
ip vrrp enable.....	157
ip vrrp preempt.....	157

ip vrrp priority.....	157
ip vrrp track.....	159
ip vrrp virtual-address.....	157
VRRP commands	
icr.....	177

W

warm restart	
IP interface priority	54
warm restart, line modules	
events that occur	
denial of new user session login attempts.....	84
preservation of active user sessions.....	84
preservation of data forwarding through chassis.....	84
prevention of configuration changes.....	84