



ERX Edge Routers Release Notes

Release 4.1.2

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

Part No. 162-00486-04 Rev. A00

Juniper
NETWORKS

Juniper Networks is registered in the U.S. Patent and Trademark Office and in other countries as a trademark of Juniper Networks, Inc. Broadband Cable Processor, ERX, ESP, E-series, G1, G10, G-series, Internet Processor, J-Protect, Juniper Your Net, JUNOS, JUNOScript, JUNOSe, M5, M10, M20, M40, M40e, M160, M-series, NMC-RX, SDX, ServiceGuard, T320, T640, T-series, UMC, and Unison are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. All specifications are subject to change without notice.

Products made or sold by Juniper Networks (including the G1 and G10 CMTSs, ERX-310, ERX-705, ERX-710, ERX-1410, ERX-1440, M5, M10, M20, M40, M40e, M160, and T320 routers, T640 routing node, and the JUNOS, SDX-300, and ServiceGuard software) 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, 6,333,650, 6,359,479, and 6,406,312.

ERX Edge Routers Release Notes, Release 4.1.2
Copyright © 2003, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writers: Brian Wesley Simmons
Editor: Fran Mues

Revision History
August 2003

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

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.

SOFTWARE LICENSE AGREEMENT¹

JUNIPER NETWORKS, INC. IS WILLING TO LICENSE THE ENCLOSED SOFTWARE AND ACCOMPANYING USER DOCUMENTATION (COLLECTIVELY, THE "PROGRAM") TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. PLEASE READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE COPYING OR USING THE ACCOMPANYING SOFTWARE OR INSTALLING THE HARDWARE UNIT WITH PRE-ENABLED SOFTWARE OR USING THE ACCOMPANYING USER DOCUMENTATION.

BY USING THE ACCOMPANYING SOFTWARE OR INSTALLING THE HARDWARE UNIT WITH PRE-ENABLED SOFTWARE, YOU AGREE TO BE BOUND BY THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. IF YOU DO NOT AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE AGREEMENT, JUNIPER NETWORKS IS UNWILLING TO LICENSE THE PROGRAM TO YOU, IN WHICH EVENT YOU SHOULD PROMPTLY WITHIN TEN (10) DAYS FROM SHIPMENT RETURN THE UNUSED SOFTWARE, USER DOCUMENTATION, AND RELATED EQUIPMENT AND HARDWARE TO THE PLACE OF PURCHASE AND YOU WILL RECEIVE A FULL REFUND OF YOUR LICENSE FEE. THIS LICENSE AGREEMENT REPRESENTS THE ENTIRE AGREEMENT CONCERNING THE PROGRAM BETWEEN YOU AND JUNIPER NETWORKS, AND IT SUPERSEDES ANY PRIOR PROPOSAL, REPRESENTATION OR UNDERSTANDING BETWEEN THE PARTIES.

1. License Grant. Juniper Networks, Inc. ("Juniper Networks") and its suppliers and licensors hereby grant to you and you hereby accept a nonexclusive, personal and nontransferable license to use the computer software and/or hardware unit with pre-enabled software, including all patches, error corrections, updates, and revisions thereto in machine-readable, object code form only (the "Software"), and the accompanying User Documentation on the Juniper Networks product owned by you and only as authorized in this License Agreement. You may make one (1) archival copy of the Software for backup purposes provided you affix to such copy all copyright, confidentiality, and proprietary notices that appear on the original. Except as authorized under this paragraph, no copies of the Program or any portions thereof may be made, in whole or in part, by you or any person under your authority or control.

The Software and User Documentation are protected under copyright laws. The title to Software and User Documentation shall remain solely with Juniper Networks and its suppliers.

Except as authorized above, you shall not: copy, in whole or in part, the Software or the related User Documentation; modify, reverse assemble, reverse compile, or otherwise translate, disassemble, or obtain source code for the Software or User Documentation, in whole or in part, or permit a third party to do so; rent, lease, distribute, sell, or create derivative works of the Software; pledge, lease, rent, sublicense or share its rights under this License Agreement; or, without Juniper Networks' prior written consent, assign or transfer its rights hereunder.

2. Juniper Networks' Rights. You agree that the Software, including the User Documentation, embodies Juniper Networks and its suppliers' and licensors' confidential and proprietary intellectual property protected under U.S. copyright law and you will use your best efforts to maintain their confidentiality. You further acknowledge and agree that Juniper Networks or its suppliers and licensors own all right, title, and interest in and to the Software, including all intellectual property rights therein. You shall take no action inconsistent with Juniper Networks or its suppliers' ownership of such Software. You shall not sublicense, assign, or otherwise disclose to any third party the Software or any information about the operation, design, performance, or implementation of the Software and User Documentation without prior written consent of Juniper Networks. You agree to implement reasonable security measures to protect such confidential and proprietary information and copyrighted material. This License Agreement does not convey to you an interest in or to the Program, but only the limited right of use revocable in accordance with the terms of this License Agreement.

3. License Fees. The license fees paid by you are paid in consideration of the license granted under this License Agreement.

4. Term. This license is effective upon opening of the package(s) or use of the hardware containing the Software, and shall continue until terminated. You may terminate this License at any time by returning the Software, including any User Documentation, and all copies or portions thereof to Juniper Networks. This License will terminate immediately without notice from Juniper Networks if you breach any term or provision of this License. Upon such termination by Juniper Networks, you

1. If you and Juniper Networks, Inc., have executed another license agreement for the Program which is now in effect, then such agreement ("Negotiated Agreement") shall supersede this Software License Agreement and shall exclusively govern the use and license terms of the Program.

must return the Software, including any User Documentation, and all copies or portions thereof to Juniper Networks. Termination of this License Agreement shall not prejudice Juniper Networks' rights to damages or other available remedy.

5. Limited Software Warranty: Juniper Networks warrants, for your benefit alone, that for a period of ninety (90) days from the date of shipment from Juniper Networks that the Software substantially conforms to its published specifications.

The limited warranty extends only to you as the original licensee. Your exclusive remedy and the entire liability of Juniper Networks and its suppliers under this limited warranty will be, at Juniper Networks' option, repair or replacement of the Software, or refund of the amounts paid by you under this License Agreement. You agree that this is your sole and exclusive remedy for breach by Juniper Networks, its suppliers or its licensors of any warranties made under this License Agreement.

In no event does Juniper Networks warrant that the Software is error free or that you will be able to operate the Software without problems or interruptions. Juniper Networks does not warrant: 1) that the functions contained in the software will meet your requirements; 2) that the Software will operate in the hardware or software combination that you may select; 3) that the operation of the Software will be uninterrupted or error free; or 4) that all defects in the operation of the Software will be corrected.

This warranty does not apply if the product: 1) has been altered, except by Juniper Networks; 2) has not been installed, operated, repaired, or maintained in accordance with instruction supplied by Juniper Networks; or 3) has been subjected to or damaged by improper environment, abuse, misuse, accident, or negligence.

EXCEPT FOR THE WARRANTIES SET FORTH ABOVE, THE SOFTWARE IS LICENSED "AS IS," AND JUNIPER NETWORKS DISCLAIMS ANY AND ALL OTHER REPRESENTATIONS, CONDITIONS, AND WARRANTIES, WHETHER EXPRESS, IMPLIED, OR STATUTORY, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR ANY WARRANTIES FOR NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. ANY AND ALL SUCH WARRANTIES ARE HEREBY EXCLUDED TO THE EXTENT ALLOWED BY APPLICABLE LAW. JUNIPER NETWORKS' SUPPLIERS AND LICENSORS DO NOT MAKE OR PASS ON TO YOU OR ANY THIRD PARTY ANY EXPRESS, IMPLIED, OR STATUTORY WARRANTY OR REPRESENTATION, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR ANY WARRANTIES FOR NON-INFRINGEMENT.

6. Proprietary Rights Indemnification. Juniper Networks shall at its expense defend you against and, subject to the limitations set forth elsewhere herein, pay all costs and damages made in settlement or awarded against you resulting from a claim that the Program as supplied by Juniper Networks infringes a United States copyright or a United States patent, or misappropriates a United States trade secret, provided that you: (a) provide prompt written notice of any such claim, (b) allow Juniper Networks to direct the defense and settlement of the claim, and (c) provide Juniper Networks with the authority, information, and assistance that Juniper Networks reasonably deems necessary for the defense and settlement of the claim. You shall not consent to any judgment or decree or do any other act in compromise of any such claim without first obtaining Juniper Networks' written consent. In any action based on such a claim, Juniper Networks may, at its sole option, either: (1) obtain for you the right to continue using the Program, (2) replace or modify the Program to avoid the claim, or (3) if neither (1) nor (2) can reasonably be effected by Juniper Networks, terminate the license granted hereunder and give you a pro rata refund of the license fee paid for such Program, calculated on the basis of straight-line depreciation over a five-year useful life. Notwithstanding the preceding sentence, Juniper Networks will have no liability for any infringement or misappropriation claim of any kind if such claim is based on: (i) the use of other than the current unaltered release of the Program and Juniper Networks has provided or offers to provide such release to you for its then current license fee, or (ii) use or combination of the Program with programs or data not supplied or approved by Juniper Networks if such use or combination caused the claim.

7. Limitation of Liability. IN NO EVENT WILL JUNIPER NETWORKS OR ITS SUPPLIERS OR LICENSORS BE LIABLE FOR ANY COST FOR SUBSTITUTE PROCUREMENT; SPECIAL, INDIRECT, INCIDENTAL, PUNITIVE, EXEMPLARY, OR CONSEQUENTIAL DAMAGES; OR ANY DAMAGES RESULTING FROM INACCURATE OR LOST DATA OR LOSS OF USE OR PROFITS ARISING OUT OF OR IN CONNECTION WITH THE PERFORMANCE OF THE SOFTWARE, EVEN IF JUNIPER NETWORKS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. Juniper Networks' cumulative liability to you or any other party for any loss or damages resulting from any claims, demands, or actions arising out of or relating to this License Agreement shall not exceed the total fees paid to Juniper Networks for the Software.

8. Export Control. Software, including technical data, is subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to

export or import regulations in other countries. You agree to comply strictly with all such regulations and acknowledge that you have the responsibility to obtain licenses to export, re-export, or import Software.

9. Government Licensees: If any Software or associated documentation is acquired by or on behalf of a unit or agency of the United States government, the government agrees that such Software or documentation is a "commercial item" as that term is defined in 48 C.F.R. 2.101, consisting of "commercial computer software" or "commercial computer software documentation" as such terms are used in 48 C.F.R. 12.212 of the Federal Acquisition Regulations and its successors and 48 C.F.R. 227.7202-1 through 227.7202-4 of the DoD FAR Supplement and its successors. The use, duplication, or disclosure by the United States government of technical, data, computer software and documentation is subject to the restrictions set forth in FAR section 12.212(a), FAR section 52.227-14(g)(2), FAR section 52.227-19, DFARS section 252.227-7015(b), DFARS section 227.7202-1(a), and DFARS section 227.7202-3(a), as applicable. All United States government end users acquire the Software with only the rights set forth in this License Agreement.

10. General: This License shall be governed by and construed in accordance with the laws of the Commonwealth of Massachusetts, United States of America, as if performed wholly within the state and without giving effect to the principles of conflict of law. Any dispute arising out of this Agreement shall be referred to an arbitration proceeding in Boston, Massachusetts, in accordance with the commercial arbitration rules of the American Arbitration Association (the "AAA"). If the parties cannot agree upon an arbitrator, arbitration shall be conducted by a neutral arbitrator selected by the AAA who is knowledgeable in electronics equipment manufacturing and software licensing. The parties shall share the procedural costs of arbitration equally, and each party shall pay its own attorneys' fees and other costs and expenses associated with the arbitration, unless the arbitrator decides otherwise. The arbitrator's award shall be in writing and shall include a statement of reasons, but the arbitrator shall not be permitted to award punitive or indirect damages. The arbitrator's decision and award shall be final and binding and may be entered in any court having jurisdiction. The terms of this section shall not prevent any party from seeking injunctive relief in any court of competent jurisdiction in order to protect its proprietary and confidential information. If any term or provision hereof is found to be void or unenforceable by a court of competent jurisdiction, the remaining provisions of this License Agreement shall remain in full force and effect. This License Agreement constitutes the entire agreement between the parties with respect to the use of the Software and User Documentation and supersedes any and all prior oral or written agreements, discussions, negotiations, commitments, or understandings. No amendment, modification, or waiver of any provision of this License Agreement will be valid unless in writing and signed by the authorized representative of the party against which such amendment, modification, or waiver is sought to be enforced. The waiver by either party of any default or breach of this License Agreement shall not constitute a waiver of any other or subsequent default or breach. This License Agreement shall be binding upon the parties and their respective successors and permitted assigns.

Should you have any questions about this agreement, please contact:

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089
Attn: Contracts Administrator

Contents

Release Installation	1
Erroneous Error Message	1
Upgrading from a 3.x Release to a 4.x Release	2
Moving Line Modules Between Releases	4
SRP Module Memory Requirements	4
Hardware and Software Compatibility	4
Release Overview	4
Before You Start	4
Contacting Customer Service	6
Release Highlights	6
HDLC	6
Release Software Protocols	7
Core Routing Stack	7
Routing Protocols	7
Multiprotocol Label Switching (MPLS)	7
Layer 2 Protocols	7
Security Protocols	8
Simple Network Management Protocol (SNMP)	8
Known System Behavior	8
ATM	8
BGP	9
B-RAS	9
CLI	9
Ethernet	10
Hardware	10
IP	10
L2TP	11
Layer 2 over MPLS	11
MLPPP	11
MPLS	12
Policy Management	12
QoS	12
System	13

Known Problems and Limitations	14
ATM	14
B-RAS	14
CLI	14
Documentation	15
Forwarding	16
HDLC	16
IP	16
MLPPP	17
MPLS	17
NTP	18
OSPF	18
QoS	18
System	19
Resolved Known Problems	19
AAA	19
ATM	20
BGP	20
CLI	20
Core DUMP	21
DHCP Relay	21
Dynamic Interfaces	21
Ethernet	21
Firmware	21
Forwarding	21
FTP Server	22
Hardware	22
IP	22
IP Multicasting	23
IS-IS	23
L2TP	23
MPLS	24
OSPF	25
PIM-SM	25
Policy Management	25
PPP	26
PPPoE	26
QoS	26
RADIUS	26
RIP	26
Routing Policy	26
SDX	27
SNMP	27
SONET	27
System	28

ERX 4.1.2 Release Notes

Release Installation

Complete procedures for installing the system software are available in the *ERX Installation and User Guide, Appendix E, Installing ERX System Software*. Before upgrading to a new version of software, save your system's running configuration to a .cnf file or .scr file. If you subsequently need to downgrade for any reason, you can restore the earlier software version.



Note: *When you upgrade the software on a system that has a large number of interfaces configured, the system may appear to be unresponsive for several minutes. This condition is normal, and you should allow the process to continue uninterrupted.*

Erroneous Error Message

Most software releases lower than 4.0.0 do not recognize the system version. Consequently, when upgrading from systems running those software versions to 4.0.0 or higher, the system will display an error message when you issue the command to boot with the 4.x release:

```
host1(config)#boot system erx_4-0-0.rel
WARNING 06/21/2002 14:47:34 os: Current system version
string invalid: 4.0.0 5
WARNING 06/21/2002 14:47:35 os: Current system version
string invalid: 4.0.0
```

You can ignore the messages as the system will boot properly. The system will also display the messages when you issue the **dir** command. If you downgrade from 4.x to a lower software version, issuing the **dir** command after booting will also display the error messages.

Upgrading from a 3.x Release to a 4.x Release

You must perform a factory reset if you want to upgrade from a release lower than 4.0.0 to Release 4.0.0 or higher. To retain your system configuration, perform the following steps:

- 1 Set the terminal length of the console to 0.
- 2 While running a release lower than 4.0.0 on a configured system, use the **show configuration** command to create a configuration script by redirecting the output to a file before you install 4.0.0 or higher software.
- 3 Verify that the script was created.
- 4 Install the 4.0.0 or higher software.
- 5 Run the configuration script after the new software has been installed and all line modules are online.

The following screen output illustrates the process.

```

host1#term length 0
host1#show configuration &> 3-2-3.scr
  -- show config output --
  ...
  ...
host1#dir

```

file	size	unshared size	date (UTC)	in use
config.cnf	531755	531755	06/27/2002 13:02:18	
reboot.hty	1536	1536	06/27/2002 13:43:40	
3-2-3.rel	69840063	69840063	06/25/2002 11:15:30	!
3-2-3.scr	149327	149327	06/27/2002 14:40:28	
erx_4-0-0.rel	87104081	87104081	06/27/2002 15:52:36	

```

host1#config t
Enter configuration commands, one per line.  End with ^Z.
host1(config)#boot config factory-defaults
WARNING: Execution of this command will cause all configuration settings to
  revert to factory defaults upon the next system reboot.
Proceed with 'boot config factory-defaults'? [confirm]

host1(config)#boot system erx_4-0-0.rel
Please wait.....
WARNING: It is recommended that you copy the current running-configuration
  to a file prior to running with a different release of software.
host1(config)#exit

```

```
host1#reload
WARNING: Execution of this command will cause the system to reboot.
Proceed with reload? [confirm]
Reload operation commencing, please wait...
  -- Reload/boot messages --
  ...
  ...
Logged in on console 0.
Copyright (c) 1998-2002 Juniper Networks, Inc. All rights reserved.
```

```
host1>enable
```

```
host1#dir
```

file	size	unshared size	date (UTC)	in use
config.cnf	531755	531755	06/27/2002 13:02:18	
reboot.hty	1536	1536	06/27/2002 13:43:40	
3-2-3.rel	69840063	69840063	06/25/2002 11:15:30	
3-2-3.scr	149327	149327	06/27/2002 14:40:28	
erx_4-0-0.rel	87104081	87104081	06/27/2002 15:52:36	!

```
* File is not stored in main file system and is not included in capacity figures
```

```
host1#show version
```

```
Copyright (c) 1998-2002 Juniper Networks, Inc. All rights reserved.
System Release: erx_4-0-0.rel
      Version: 4.0.0 release-0.0 (June 23, 2002 20:03)
System running for: 0 days, 0 hours, 4 minutes, 12 seconds
      (since WED JUN 27 2002 16:03:21 UTC)
```

slot	state	type	admin	spare	running release
0	online	SRP-10Ge	enabled	---	erx_4-0-0.rel
1	---	---	---	---	---
2	online	TUNNEL-SERVER	enabled	---	erx_4.0.0.rel
3	online	OC3d	enabled	---	erx_4-0-0.rel
4	online	GE(P1)	enabled	---	erx_4-0-0.rel
5	online	COC3-4	enabled	---	erx_4-0-0.rel
6	---	---	---	---	---

```
host1#conf f 3-2-3.scr
```

```
Proceed with configure? [confirm]
host1#
```

Moving Line Modules Between Releases

The ERX-1440 router employs a 40-Gbps SRP module and a new midplane. Release 3.3.2 was the first software release to support the 40-Gbps SRP module and midplane. Before you can transfer a line module from an ERX-705, ERX-710, or ERX-1400 router to an ERX-1440 router, you must first load Release 3.3.2 or a higher release onto the current router, and then reboot the router to load the release onto the line modules. If you then move any of those line modules to an ERX-1440 router, that router is able to recognize the line module.

If you move a line module from an ERX-1440 router to an ERX-705, ERX-710, or ERX-1400 router, the module will load properly in the new router regardless of the release.

SRP Module Memory Requirements

You must have at least 512 MB of memory on the SRP module to run the current software release.

Hardware and Software Compatibility

For important information on the compatibility of hardware and software, allowable combinations of line modules, and memory requirements, see *ERX Release Notes 4.1.1, Appendix B, ERX Modules*.

Release Overview

These Release Notes cover Release 4.1.2 of the system software for the ERX-700 and ERX-1400 series edge routers. If the information in these Release Notes differs from the information found in the published documentation set, follow these Release Notes.

Topic	Page
Release Highlights	6
Release Software Protocols	7
Known System Behavior	8
Known Problems and Limitations	14
Resolved Known Problems	19
<i>Appendix A, System Maximums</i>	A-1

Before You Start

These *Release Notes* include information on the changes between Releases 4.1.1 and 4.1.2. Before you use your new software, read these

Release Notes in their entirety, especially the section *Known Problems and Limitations*. You need the following documentation to fully understand all the features available in Release 4.1.2:

- These 4.1.2 Release Notes, which describe features available in Release 4.1.2.
- The 4.1.1 Release Notes, which describe features available in Release 4.1.1.
- The 4.1.0 Release Notes, which describe features available in Release 4.1.0.
- The 4.1.x ERX system documentation set, which provides detailed information on features available in Release 4.1.0.

If the information in your current Release Notes differs from the information found in the other documentation sources, follow the Release Notes.

The 4.1.x ERX documentation set consists of several manuals and is available in electronic (PDF) or printed format. Refer to the following table to help you decide which document to use.

If you want to . . .	Go to . . .
Install the system	<i>ERX Installation and User Guide</i>
Get up and running quickly	<i>ERX Installation Quick Start poster</i>
Configure the system	<i>ERX System Basics Configuration Guide</i>
Configure physical layer interfaces	<i>ERX Physical and Link Layers Configuration Guide, Part 1</i>
Configure link layer interfaces	<i>ERX Physical and Link Layers Configuration Guide, Part 2</i>
Configure routing management	<i>ERX Routing Protocols Configuration Guide, Vol. 1</i>
Configure routing protocols	<i>ERX Routing Protocols Configuration Guide, Vol. 1, Part 1</i> <i>ERX Routing Protocols Configuration Guide, Vol. 2</i>
Configure policy management and quality of service (QoS)	<i>ERX Policy and QoS Configuration Guide</i>
Configure remote access	<i>ERX Broadband Access Configuration Guide</i>
Get specific information about commands	<i>ERX Command Reference Guide</i>
Get a high-level overview of the ERX system and its architecture	<i>ERX Product Overview Guide</i>

Contacting Customer Service

You can contact Juniper Networks Customer Service for the ERX product line in the following ways:

- Within the United States, call 1-888-314-JTAC
- Outside the United States, call 408-745-9500
- Send e-mail to support@juniper.net

Release Highlights

Release 4.1.2 is a maintenance release, but does include the following software feature.

HDLC

- New MRU and MTU Maximums for HDLC Interfaces

On the CE1, CT1, cOCx/STMx, CT3, E3-FRAME, HSSI, POS, T3-FRAME, and X.21/V.35 line modules, the maximum allowable size in bytes of the MRU and MTU for HDLC interfaces has been changed to 10,000 minus the CRC size.

The CRC size can be set to 16 (2 bytes), 32 (4 bytes), or none. The total value of the MRU/MTU size and CRC size combined cannot exceed 10,000 bytes on any line module.

If you set the MRU/MTU size to a value greater than 10,000 minus the CRC, the router issues a warning and automatically sets the MRU/MTU size to one of the following values:

- > 10,000 if the CRC is set to none
- > 9998 if the CRC is set to 16
- > 9996 if the CRC is set to 32

If you attempt to increase the CRC to a value that exceeds the 10,000-byte maximum after the MRU/MTU size is set, the **crc** command will fail. For example, if the MTU for a POS module is set to 9998, attempting to increase the CRC value from 16 to 32 would result in a combined value that exceeds 10,000 bytes, and will therefore cause the **crc** command to fail.

Release Software Protocols

The following list identifies the major software protocols supported in this release.

Core Routing Stack

- Internet Protocol (IP)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)

Routing Protocols

- Border Gateway Protocol (BGP-4)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Group Membership Protocol (IGMP)
- Protocol Independent Multicast Protocol (PIM), including PIM Dense Mode, PIM Sparse Mode, and PIM Dense-Sparse Mode
- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP) version 2
- Virtual Router Redundancy Protocol (VRRP)

Multiprotocol Label Switching (MPLS)

- Label Distribution Protocol (LDP)
- Constraint-based Routing Label Distribution Protocol (CR-LDP)
- Resource ReSerVation Protocol - Traffic Engineering Extensions (RSVP-TE)

Layer 2 Protocols

- Asynchronous Transfer Mode (ATM)
- Cisco HDLC
- Ethernet
- Frame Relay (FR)
- Layer 2 Tunneling Protocol (L2TP)
- Multilink Frame Relay (MLFR)

- Multilink Point-to-Point Protocol (MLPPP)
- Point-to-Point Protocol (PPP)
- PPP over Ethernet (PPPoE)

Security Protocols

- IP Encapsulating Security Payload (ESP)
- IP Authentication Header (AH)
- Internet Security Association and Key Management Protocol (ISAKMP)
- Internet Key Exchange (IKE)

Simple Network Management Protocol (SNMP)



Note: For detailed information on any protocol, see the Configuration Guides.

Known System Behavior

This section briefly describes ERX system behavior and related issues. In some cases the behavior differs from non-ERX implementations; in others the behavior is included to emphasize how the system works.

ATM

- The ERX system will decrease the maximum number of subscribers allowed when the memory limit is approached on ATM modules with 128 MB of memory. This might happen in either of the following circumstances:
 - > In oversubscription configurations when 16k ATM circuits are configured for dynamic subscribers, 8k IP/PPP/PPPoE/ATM interfaces are active, and another 8k subscribers continue their attempts to log in.
 - > In endurance tests with 8k PPP/PPPoE dynamic interfaces that have elaborate policies attached to them.

Work-around: Upgrade the module to 256 MB.

BGP

- If you specify an interface when you issue the **neighbor update-source** command and the interface is later removed, then this command is also removed from the system configuration.
- The following message might be displayed under certain conditions:

```
bgpConnections (default,0.0.0.0): Error accepting inbound  
TCP connection
```

The message is generated when an unconfigured peer attempts to establish a TCP session with an ERX system and a valid route to the source address of the peer is absent from the router's routing table.

If a valid route exists in the routing table, the following message is displayed when an unconfigured peer attempts to establish a TCP session with an ERX system; X.X.X.X is the source address of the unconfigured peer:

```
NOTICE 08/29/2001 16:50:11 bgpConnections (default,X.X.X.X):  
Inbound connection refused - no peer X.X.X.X configured in  
core
```

B-RAS

- The OC3 module (non-ASIC, dual-port) does not support rate shaping of egress traffic for VP tunnels. It does support rate limiting, where traffic exceeding the rate is dropped. When traffic is dropped, the link bandwidth corresponding to the discarded packets is not reallocated to other tunnels or individual VCs, resulting in unused bandwidth.
- Pool groups are not supported; although the **ip local pool group** command appears in the CLI, it is not supported.

CLI

- The **show config** command normally takes a long time to finish for extremely large configurations. If you specify a search string (via the **begin**, **exclude**, or **include** options) with the command for a string that is not present in the configuration, then the CLI session appears to be busy for a prolonged period. The CLI filtering feature for **show** commands does not speed up execution of the command.
- If you remove a standby SRP module and fail to issue the **slot erase** command, then the ERX system cannot guarantee that the SRP modules were synchronized. In this situation, the ERX system

properly will not execute a simple **reload** command. To reload the router you must now do either of the following: [Defect ID 42219]

- > Issue the **reload force** command.
- > Issue the **slot erase** command followed by the **reload** command.

Ethernet

- When counting bits per second on a Fast Ethernet or Gigabit Ethernet interface, the ERX system includes 12 bytes for interpacket gap, 7 bytes for preamble, and 1 byte for start frame delimiter, for a total of 20 bytes (160 bits) per packet more than some non-ERX systems. This value therefore shows the total bandwidth utilization on the interface, including both data and overhead.

Hardware

- There is a very small chance that some line modules could have an improperly modified keying block that prevents the module from proper seating in the top slot of an older ERX-700 chassis.
Work-around: Remove the keying block to insert the module in the top slot, or insert the module in a different slot.
- Because the FE-2 line module connects directly to its I/O module rather than to a midplane, installing one module can unseat the other module.
Work-around: Ensure that the screws fastening the I/O module to the chassis are tight to prevent separation between the modules and a resulting poor connection.
- Use the Ethernet port on the SRP I/O module only as a system management port. If you must route Ethernet on a 10/100BaseT port, you must do so on an FE-2 or FE-8 I/O module.

IP

- ASIC line modules can accommodate routing tables of up to 2M entries. Non-ASIC line modules are limited to 100K entries (although more routes could still work because a routing table of up to 1M entries could be compressed to 100k or fewer nonconsecutive routes). See *ERX Installation and User Guide, Appendix B, Module Specifications*, for information on which modules are ASIC and which are non-ASIC.

- When you issue the **show ip forwarding-table** command for a particular slot, it is normal and appropriate behavior if the Status field indicates Valid while the Load Errors field is increasing daily for that VR. The Load Errors field records any failed routing table distribution attempt as an error. Attempts can fail for many reasons during normal operation; a failed attempt does not necessarily indicate a problem. It is normal to see many Load Errors per day. If the Status field does not indicate Valid, then the routing table distribution has failed constantly for that VR, and a real problem exists.

L2TP

- If you create an L2TP destination profile *profileName*, establish tunnels with the profile, and then remove the profile, you cannot subsequently create another destination profile using that same *profileName* until all the tunnels drain from the previous instance of this destination profile. If you do not wait, the ERX system displays a message similar to the following:

```
l2tp: Discarding incoming sccrq from vr default, remote
address 192.168.100.1 - no destination profile.
```

If you do not want to wait for the tunnels to drain, you can simply use a different name for the destination profile. [Defect ID 32973]

Layer 2 over MPLS

- You can enter the **router interface tunnel mpls:baseTunnel** command in any virtual router or VRF context. However, the configuration takes place in the virtual router in which the base tunnel resides.

You must enter the **mpls-relay remoteAddress** command in the virtual router where the *remoteAddress* can be reached; that is, in the virtual router providing core connections. This means, for example, that you cannot enter the command in a VRF.

MLPPP

- An ERX system deployed as an LNS where the tunnel sessions are using MLPPP can have only a single Tunnel Service module installed. Having more than one TSM installed in this configuration leads to incorrect operation of the ERX system.
- The **ppp chap-challenge-length** command is not currently supported for MLPPP. [Defect ID 32040]

MPLS

- The ERX system does not support multiple FEC elements in a single label mapping message, which is the default behavior for Juniper Networks M-series routers. When communicating with an M-series router that uses the default setting, an ERX system might display the following error message:

```
ERROR 04/09/2003 19:46:00 mplsGeneral (default):
  LDPGetFecElemsFromTlv:too many FEC Elements in TLV.
  Maximum number supported is 1.
```

Work-around: Use the **set protocol ldp deaggregate** command on the M-series router to specify that the router not include multiple FEC elements in a single label mapping.

- Do not enable OSPF on any unidirectional interfaces (such as an MPLS tunnel) because it will never be able to form an adjacency. If the range specified by the **network area** command includes an address on an interface that is being referred to by unnumbered interfaces, all of the unnumbered interfaces will begin trying to form adjacencies. If this behavior is not intended, you must reevaluate the interface assignment or the range specified by the command.

Policy Management

- Traffic shaping is not supported on non-ASIC modules. For information on which modules are ASIC or non-ASIC, see *ERX Installation and User Guide, Appendix B, Module Specifications*.

QoS

- A dynamic IP interface can have a QoS profile attached by RADIUS. For example, if configured by RADIUS, the **show ip interface** command might show the following:

```
Attached QoS profile: Strict-qos
```

However, if the profile is configured statically, the QoS profile would be attached to the 1483 subinterface, and the attachment would be displayed by the **show atm subinterface** command rather than **show ip interface**.

- In Release 4.0.0 and higher releases, ATM VCs are not shaped to the configured rate if you use the **atm-vc node scheduler-profile** command to configure VC shaping and your QoS configuration involves ATM VPs shaped in the SAR.

This command does shape VCs correctly if you have configured the SAR for per-port queuing via the **qos-mode-port** command.

However, per-port queuing is recommended mostly for configurations with strict-priority scheduling. You cannot issue the **qos-mode-port** command if you have configured the SAR for VP tunnel shaping, because the command disables all ATM/SAR shaping configurations.

Work-around: Configure shaping on the IP interface, rather than on the VC, via the **ip node scheduler-profile** command. The **ip node scheduler-profile** command is added to the new QoS profile and any existing **atm-vc node** commands are left in place. The **ip node scheduler-profile** command is also compatible with the default mode of operation of the SAR. Using the **ip node scheduler-profile** command works well in configurations with one IP interface per VC, which is the case except for when multiple PPPoE clients are stacked above a single VC.

System

- In a router with a redundancy group that does not span quadrants (for example, a three-slot redundancy group that spans slots 0, 1, and 2 in a 14-slot chassis), the potential bandwidth of the redundant module is erroneously included in the quadrant bandwidth calculation. The **show utilization** command might indicate that the bandwidth is exceeded for modules in that group. [Defect ID 31034]
- When you copy the running configuration to NVS, the ERX system checks whether it has available space equal to at least twice the size of .cnf file. If the space is insufficient, you cannot complete the copy. [Defect ID 40655]

Work-around: Make sufficient space on the NVS by deleting release or .cnf files.

- If you upgrade the system software on a system with redundant SRP modules, the secondary SRP module does not run the new software until it reboots. If, before it is booted, you issue the **srp switch** command or the primary SRP module fails, the secondary SRP module will run with the old release when it takes control.
- The **show version** command displays how long the system (meaning the chassis) has been running; it does not reflect the uptime of a particular SRP module.

Known Problems and Limitations

This section identifies the known problems and limitations in this release. For more information on known problems that were discovered at customer sites, you can go to the Customer Support Center Web site at <https://www.juniper.net/support/csc/kb.html>, log in to the ERX Knowledge Base, enter the defect number in the Search By Keyword field, and click Search.

ATM

- When you reload an ATM line module that is configured with NBMA circuits as passive OSPF interfaces and has established OSPF adjacencies and IBGP peers (configured on Gigabit Ethernet interfaces), the transmission of OSPF hello packets may be affected until all the NBMA interfaces have initialized. [Defect ID 46157]

Work-around: Either remove the passive OSPF interface statements on the NBMA interfaces or statically configure the OSPF cost on the NBMA interfaces.

- If you configure 1:5 or 1:4 redundancy groups with ATM modules and more than 3 modules are active in the redundancy group, the ERX system might exhibit poor forwarding and control performance when an active line module fails over to a standby module. [Defect ID 43257]

B-RAS

- If the ERX system is under a heavy load, the **show profile** command might take longer than usual to execute. [Defect ID 41738]

Work-around: You can either delay examination of profiles until the system is less busy or you can save a copy of the profile to a text file off the ERX system.

CLI

- The **show subsystem** and **exclude-subsystem** commands show the TSM as an OC12 module. [Defect ID 36661]
- You cannot use an underscore character (`_`) in an MPLS tunnel name unless you enclose the entire name in quotes. [Defect ID 31291]

Documentation

- Units Incorrect for **show isis spf-log** Command

ERX Routing Protocols Configuration Guide, Vol. 1, Chapter 10, Configuring IS-IS, incorrectly states that the Duration field displays number of milliseconds to complete the SPF run. The field instead displays the value in seconds; for example 0.160.

- 4XOC3/STM1 I/O modules with APS/MSP redundancy are unavailable.

Some versions of the ERX documentation incorrectly state that the 4XOC3 APS MULTIMODE and 4XOC3 APS SINGLE MODE I/O modules, which pair with the OC_x/STM_x ATM and OC_x/STM_x POS line modules to provide four OC3/STM1 ports with APS/MSP redundancy, are available. These modules are not currently available.

- The *ERX Command Reference Guide* incorrectly describes the **aaa timeout** command:
 - > The correct value for the maximum session timeout that you can specify for B-RAS PPP user is 1814400 seconds (21 days), not 2678400 (31 days).
 - > If you enter a timeout value that is not a multiple of 60 seconds; the system accepts the value as is and does not round it to the nearest minute.
 - > If you enter a value that is out of range, the CLI displays a warning message; the values are not rounded to be within the accepted range.
 - > Although the CLI does not enforce the rounding rules stated in the guide for values you enter, the timeout values can be provided by RADIUS. The range is not enforceable on such values. PPP and L2TP will round these timeout values from RADIUS as follows:
 - If the timeout is less than the minimum, it is rounded up to the minimum.
 - If the timeout is greater than the maximum, it is rounded down to the maximum.
 - If the timeout is not a multiple of 1 minute, it is rounded to the nearest minute.

Forwarding

- The **show ip interface detail atm** command displays a discrepancy between the dropped byte count and the dropped packet count. [Defect ID 42501]

Work-around: Use the dropped packet count for valid information; the dropped byte data is incorrect.

HDLC

- If you attempt to increase the CRC to a value that exceeds the 10,000-byte maximum after the MRU/MTU size is set, the **crc** command will fail. For example, if the MTU for a POS module is set to 9998, attempting to increase the CRC value from 16 to 32 would result in a combined value that exceeds 10,000 bytes, and will therefore cause the **crc** command to fail. [Defect ID 50101]

Work-around: Set the MRU or MTU to a value of 10,000 minus the new CRC value before you change the CRC size.

IP

- If you issue the **show arp** command when there is a large ARP table, OC3 line module CPU utilization can spike to 80-90%. [Defect ID 33126/43091]

Work-around: Issue the **show arp** command for a particular interface rather than for the entire module.

- The FE-2 module counters might display incorrectly for the In Fabric Dropped Packets statistic when tunneled traffic is forwarded across the module. [Defect ID 27308]
- The **ip route permanent** command does not work properly. [Defect ID 34303]

Work-around: You can issue the **ip alwaysup** command to prevent the route from being removed from the IP routing table after the interface is shut down.

- IP interface statistics become inconsistent when a slot is reset, because some traffic (such as control traffic) might be destined for the SRP module and is therefore counted elsewhere. [Defect ID 26697/2445]
- If you have enabled ipInterface logging at a priority of debug, the acknowledgement that an interface has been deleted from the line modules can return to the SRP module after the layers beneath IP have deleted their interfaces. Consequently, the original name of the

interface cannot be resolved or displayed in the log, and the system instead displays the `ifIndex` of the IP interface. [Defect ID 32624]

Work-around: This behavior has no functional effect other than that the log is misleading. However, previous log events will have logged that the interface deletion was beginning.

- A ping across an ECMP link with the DF bit set and the data size larger than the MTU size of one of the ECMP paths may fail because the ICMP CRC of the ping is corrupted. [Defect ID 41795]

Work-around: Configure the same MTU size on all of your ECMP paths and do not set the DF bit.

MLPPP

- You cannot configure the CHAP challenge length for static MLPPP interfaces. We recommend against modifying the default CHAP challenge length. [Defect ID 32040]
- Multilink PPP does not detect illegal bundle members.

Work-around: Display the local and peer endpoint discriminators via the `show mlppp interface full` command to help debug the problem. [Defect ID 3012]

MPLS

- If you explicitly configure MPLS on an interface that is an inter-AS BGP/MPLS VPN interface, the inter-AS VPN will not work.

Work-around: Do not explicitly configure MPLS on an inter-AS interface. The ERX implementation of inter-AS BGP/MPLS VPNs does not require or support such a configuration. [Defect ID 47067]

- For layer 2 circuits over MPLS, the VC ID must be unique per VC type across all peers. This is in contrast to the Martini draft, which permits the VC ID to be the same for different peers.

For example, the draft would permit VC ID 10 to be used for a particular VLAN subinterface connected to a peer over the MPLS network, while VC ID 10 is also to connect another VLAN subinterface to a different peer. The ERX system requires that different VC IDs be used for different layer 2 circuits, even when they're connected to different peers.

An exception to this behavior is for locally switched layer 2 circuits where the two local circuits are cross-connected and therefore have the same VC ID. [Defect ID 47819]

- If LSPs are announced into IGPs, then the IGP routes cannot be used for multicast RPF checks, because LSPs are unidirectional.

Work-around: Configure static RPF routes with native hops if LSPs are auto-route announced to IGPs. [Defect ID 28526]

NTP

- The **ntp server** command does not support the **key** option. [Defect ID 28963]

OSPF

- If you configure OSPF on an interface and then remove it, a **show** command might indicate that OSPF is still running; this is only a display issue and does not affect system behavior. [Defect ID 35641]

QoS

- In Release 4.0.0 and higher releases, ATM VCs are not shaped to the configured rate if you use the **atm-vc node scheduler-profile** command to configure VC shaping and your QoS configuration involves ATM VPs shaped in the SAR. [Defect ID none]

This command does shape VCs correctly if you have configured the SAR for per-port queuing via the **qos-mode-port** command.

However, per-port queuing is recommended mostly for configurations with strict-priority scheduling. You cannot issue the **qos-mode-port** command if you have configured the SAR for VP tunnel shaping, because the command disables all ATM/SAR shaping configurations.

Work-around: Configure shaping on the IP interface, rather than the VC, via the **ip node scheduler-profile** command. The **ip node scheduler-profile** command is added to the new QoS profile and any existing **atm-vc node** commands are left in place. The **ip node scheduler-profile** command is also compatible with the default mode of operation of the SAR. Using the **ip node scheduler-profile** command works well in configurations with one IP interface per VC, which is the case except for when multiple PPPoE clients are stacked above a single VC.

- QoS is not yet supported for L2TP LACs. [Defect ID none]

System

- If you define an interface and then apply a description to it, you cannot then apply a profile to the interface. [Defect ID 42971/42746]

Work-around: You must use the following sequence:

- a Define the interface.
 - b Apply a profile to the interface.
 - c Apply a description to the interface.
- You cannot delete the ipInterface log after you delete the corresponding IP interface. This does not prevent you from adding filters to other interfaces nor does it prevent you from adding a filter to the same interface if you recreate it after deletion.

Work-around: Remove the filter before you remove the interface. Alternatively, if you remove the interface first, then you must remove all filters associated with all IP interfaces. [Defect ID 34842]

- In Interface Configuration mode for a major interface, the CLI displays options for protocols not supported by that interface type. [Defect ID 33307]

Resolved Known Problems

This release is based on the 4.1.1p1-3 patch release and incorporates all problem resolutions found in that release. This section lists significant known problems that have been resolved through that patch and later.

AAA

- SRP module reset type: panic task: radius1 file:radiusAuth.cc line: 6017 - when AAA duplication is turned on. [Defect ID 49692]
- SRP module reset: Panic in task: aaaServer file: osSemaphore.cc, line 51 when deleting radius server. [Defect ID 49408]
- The router will ignore a RADIUS attribute PPPoE URL when the **ip local pool** command is used to provide an IP address to a client. [Defect ID 48815]
- Unisphere still used instead of Juniper in aaaServerGeneral log messages. [Defect ID 48191]

ATM

- Certain ADSL modems fail to connect to ERX system; failure to create PPPoE interfaces. [Defect ID 48291]
- A **slot erase** command failure: % configuration purge failed. [Defect ID 49474]
- Incorrect ifOperStatus integer reported via SNMP. [Defect ID 49028]
- InARP packets are discarded on dynamic interfaces. [Defect ID 47414]
- F5 OAM loopback cell processing VCC. [Defect ID 48444]
- RFC1577 Lucent CellPipe DSL CPE sends inArp packet to ERX system with invalid source hw type/length value. [Defect ID 46706]
- SRP module resets with exception 0x300 (data access) in task: atm1483 when running macro to configure ATM subinterfaces. [Defect ID 46509]
- Line module reset: reset type: panic, task: scheduler file: 83DataService.cc line: 8260 last errno: 0. [Defect ID 46506]

BGP

- SRP module reset type: panic, task: bgp276, file: gpDestination.cc, line: 1104. [Defect ID 48899]
- BGP will always use next-hop-self for locally-originated routes. [Defect ID 49398]
- MIB walk of the bgpPeerIdentifier gets stuck in loop. [Defect ID 48588]
- SRP-40G resets with BGP panic in DirectNextHop.cc, line 285. [Defect ID 47874]
- The **neighbor update-source** command in a VRF should accept a loopback in the VRF (not a loopback in the parent VR). [Defect ID 47828]
- Import map configured in one VRF affects the other VRF. [Defect ID 48084]

CLI

- On rare occasions, if you delete a virtual router while another user is configuring that VR, the SRP module could reset. [Defect ID 42564]
- After a reload, the command log destination syslog x.x.x.x include x will no longer be present in the config. [Defect ID 49380/49411]

- Limit the name of map lists to no more than 31 characters. A longer name can cause the router to reset. [Defect ID 47196]
- SRP module resets with panic, task: telC15, file: osLog.cc, line: 4316. [Defect ID 48209]

Core DUMP

- Core dump from standby SRP module fails due to the filename is too long. [Defect ID 46146]

DHCP Relay

- In a DHCP relay dynamic interfaces environment, only one bridged user per VC is successfully installed. [Defect ID 50118]

Dynamic Interfaces

- Cannot delete profile that contains the **ip igmp** command. [Defect ID 48509/38866]

Ethernet

- Gigabit Ethernet with SFP IOAs shows no link light while the **show int gi** command reports link up. [Defect ID 48361]
- GE stops forwarding packets when a Cisco GSR is reloaded. [Defect ID 48663]

Firmware

- PPP over OC48 drops due to keepalive failure when running 64-byte packets at line rate. [Defect ID 47617]

Forwarding

- IP traffic statistics inconsistencies with ICMP redirects. [Defect ID 49427]
- Line module reset: GE, FC processor exception 0x700 (program), task: Egress FC[1], fpEgressProcessWhiteRouteTag - Mok_ERX_2. [Defect ID 49400]
- Line module fragmentation does not work after a fast reset. [Defect ID 48996]

- Need ability to **clear df-bit** because **ignore-df-bit** generates packets with both dont-fragment and more fragment bits set. [Defect ID 48563]
- Line module reset with exception 0x700 (program), task: Ingress FC[2] after upgrading. [Defect ID 45778]
- PPPoE clients cannot log into the ERX system. [Defect ID 48433]
- Line module reset with processor exception 0x300 (data access) task: Egress FC[1]. [Defect ID 48165]

FTP Server

- FTP to router resulted in file with 0 byte size. [Defect ID 46913]

Hardware

- SRP module resets with panic, arg (0xe7000008), file: statics.cc, line: 132, task: bgp. [Defect ID 47136]

IP

- CT3 module reset with panic in file: rfaceDatabase.cc line: 7696 task: scheduler due to an ECMP pathCount that was changed and then accessed resulting in uninitialized memory read. [Defect ID 49057]
- RFC 2547 forwarding is not working and dynamic IP interfaces are in 'not present' state. [Defect ID 49471]
- If you have a large B-RAS configuration (4000 PPPoA connections) on a T3 ATM module combined with a CT3/T3 I/O module (sometimes referred to as a 3-port T3 ATM line module and I/O) and the T3 ATM module reloads while you are running Telnet, the Telnet sessions may drop. [Defect ID 31618]
- Static route statements with next hop null 0 is replaced with 255.255.255.255 after upgrading from 4-1-0 to 4.1.1. [Defect ID 49377]
- SRP module reset: panic, file: iParserRouter.cc, line: 92, task: telC4 when entering **ip vrf forwarding vrfabc** command in a VRF. [Defect ID 49193]
- Route table distribution can leak ARF memory on the line module. [Defect ID 49558]
- The syntax for **show virtual-router** command is inconsistent from help to actual CLI. [Defect ID 46703]

- The router does not accept a next-hop address as part of a RADIUS returned framed IP route. [Defect ID 47968]
- No Telnet session possible when lines appear to be available. [31564]
- The **ip demux-type da-prefix** command does not appear in the **show config** command output for unnumbered interfaces. [Defect ID 46981]
- SRP module reset (ar1IpEngineRouteTable.cc, line 3045) on L2TP LAC after bouncing line modules on L2TP LNS. [Defect ID 47078]
- SRP module resets with processor exception 0x300 (data access) task: ip_DataCtrl. [Defect ID 47770]
- SRP module resets with exception 0x68616c74 (halt) task: ip_Ctrl_1 when “IP tunnel Reassembly” is enabled. [Defect ID 48763]
- SRP module resets with panic in ar1IpEngine.cc, line 10115 following OC3 POS panic in NextHopFecApi.cc, line 71. [Defect ID 47680]

IP Multicasting

- Console hangs after disabling/enabling multicast routing. [Defect ID 48156]
- SRP module resets with panic in task: pim1 file: mgmtEvents.cc (disable/enable multicast). [Defect ID 48273]

IS-IS

- The command **mpls traffic-eng router-id loopback0** disappears from the IS-IS configuration after you change the IP address of loopback 0. [Defect ID 49356]
- The ERX system and GSR do not interoperate for some IS-IS metric-style configurations. [Defect ID 48811]

L2TP

- SRP module reset: reset type: panic; task: l2tp; file: tpDestination.cc. [Defect ID 49138]
- An idle timeout acts like a session timeout when passed via ERX LAC. [Defect ID 47942]
- Debug code to dump l2tpShowDebugStats and l2tpShowDebugEvents after a failed tunnel or session. [Defect ID 48577]

MPLS

- L2TP tunnel over BGP/MPLS VPN configuration issue. [Defect ID 49544]
- TTL modification on egress when PHP is being performed. [Defect ID 48923]
- MPLS tail tunnel stops receiving traffic. [Defect ID 49357]
- SRP module reset: processor exception 0x68616c74 (halt), task: system. [Defect ID 49461]
- LDP sessions fail to stay up when processing large amounts of LDP data. [Defect ID 49444]
- LDP sessions down and will not come back up. [Defect ID 49139]
- Slow fail over time with 3000 Martini circuits in certain scenarios. [Defect ID 49064]
- Some MPLS shim interfaces stuck in Label Sent state after failover of primary link connecting PE to core. [Defect ID 48570]
- The LDP “labels accum” and “labels learned” counters not accurate in 4-1-0. [Defect ID 48621]
- SRP switch Slot 6 / task: mplsMgr1r1 reset type: exception 0x300 (data access). [Defect ID 48561]
- Panic in mplsMgr3r4/MplsAppService.cc/616 during L2 over MPLS reroute test. [Defect ID 49008]
- Ever-increasing MPLS tunnel number. [Defect ID 49103]
- LDP needs to include more information for TCP connect failure. [Defect ID 49043]
- LDP advertised label mappings for routes with null0 as next hop. [Defect ID 49096]
- Shim interfaces do not restack after changing MTU and bouncing Ethernet over MPLS interface. [Defect ID 48971]
- Under VR2 while disabling MPLS Entry: 19, Slot: 7, Task: mplsMgr4r5, File: rsFiconLsm.cc, Line: 511. [Defect ID 46877]
- SRP module resets with exception 0x300 (data access), task: mplsMgr1r1, LDPCfgGetFECInfo due to show mpls ip binding | i release. [Defect ID 47973]
- SRP module resets in task: mplsMgr1r1; file: inorInterface.cc; line: 349. [Defect ID 48403]

- LDP sessions would not successfully negotiate, BGP connection with peer down, forwarding table contains unresolved route. [Defect ID 48732]
- Line module reset: GE, FC processor exception 0x700 (program), task: Egress FC[1], PROC EXC 0x300 on Ingress and Egress routers (mplsMgr1r2) after MPLS goes down on Ingress. [Defect ID 44117]
- OSPF routes pointing to MPLS minor interfaces. [Defect ID 48357]
- Redundant links in core do not fail/switch over correctly in MPLS L2 VPN setup. [Defect ID 48315]

OSPF

- SRP module reset: panic in file: osPool.cc line: 657 task: ospf_OspfMsgQueue, when removing OSPF configuration. [Defect ID 49683]
- Summary link not advertised into another area after issuing **interface shut/no shut** commands. [Defect ID 49190]
- The router will not accept an empty opaque OSPF LSA if it is not max aged. [Defect ID 48770]
- OSPF adjacencies drop after removing multiple interfaces from OSPF. [Defect ID 48465]
- Removing a passive interface from OSPF results in an AS external LS being added to the LSDB. [Defect ID 48021]

PIM-SM

- The SRP module reset with a panic due to wrong prefix length in received PIM-SM packet. [Defect ID 48778]

Policy Management

- Functions as designed: Baseline is not supported for policies applied in profiles. [Defect ID 35945]
- Ascend data filter will have “Admin State” and “Operational State” as disabled. [Defect ID 49102]
- Policy baseline enable not stored in nonvolatile memory. [Defect ID 36037]
- Unable to upgrade from 4.1.1p0-2 to 5.0.1b0-12 & 5.1.0b1-10, resulting in system panic on policyMgr. [Defect ID 48600]

PPP

- Support operation in half-unnumbered mode even when RADIUS config requires numbered mode. [Defect ID 49575]

PPPoE

- PPP login test that takes 20 minutes on 3-2-3 takes several hours on 4-1-0. [Defect ID 46237]

QoS

- Line module reset: panic in task: scheduler; file: fc2QoSAgent.cc; line: 1485. [Defect ID 48821]
- QoS can't attach profiles to B-RAS interfaces. [Defect ID 48942]
- Memory leak in QoS when bouncing B-RAS users. [Defect ID 48828]
- TSM reset (fc2QoSAgent.cc, line 4170) after about 6,500 L2TP LNS sessions came down. [Defect ID 47950]

RADIUS

- SRP module reset: processor exception 0x300 (data access), task: radiusClient_rad. [Defect ID 49622]
- SRP module reset (exception 0x400, task: radiusClient_rad) on L2TP LNS as soon as scaling test began. [Defect ID 48350]
- SRP module resets with panic, file: osPool.cc @ line 855, task: radius1. [Defect ID 47044]

RIP

- SRP panic; task: rip_ripMsgQueue_, file: riprx.c, line: 1313. [Defect ID 49379]
- SRP module reset: panic, file: ripSummary.cc, line: 165, task: rip_ripMsgQueue. [Defect ID 49487]
- RIP hold timer in ERX RIP implementation causes network problems and breaks RFC 2453. [Defect ID 48861]

Routing Policy

- The SRP module reset with a processor exception 0x68616c74 (halt), task: system, after upgrade from Release 5-0-1p1 to 5-0-1p1-2. [Defect ID 49982]

- SRP module resets with panic in task: ip_89_init, File:ExtCommPolicy.cc in line 981. [Defect ID 47137]

SDX

- The ERX system is not sending the nasPortId speed through COPS connection of GRE tunnel interface. [Defect ID 42530]
- DHCP client intermittently gets into state in which it does not get address for 30 minutes due to internal cleanup/deletion of policy lists not taking place before the token iterator was deleted. [Defect ID 49805]
- SDX client does not contain pppLoginName or radiusClass variables for dynamic ATM authenticated bridged interface. [Defect ID 49624]
- GRE tunnel is reported inconsistently to SDX software, causing SDX to not know of user's existence. [Defect ID 49662]
- COPS RPT and DRQ missing for Token->Public transition. [Defect ID 47441]
- SDX client will hang when interface description is larger than 127 characters. [Defect ID 48417]

SNMP

- Functions as designed: Passwords are not displayed in CLI **show** command output after an SNMP user is created. [Defect ID 39693]
- If you use SNMP to configure GRE tunnels on the ERX system, you cannot completely configure all parameters related to the GRE tunnel unless you use a specific order of SNMP sets. [Defect ID 41796]
- SNMP FanStatusChange traps are not sent if the SNMP-Server host has a trapfilter of Notice. [Defect ID 48457]
- MPLS minor and major interfaces don't register properly with protocol layer registry. [Defect ID 45542]
- SRP module resets in mpRouterAgent.cc line 5628 during SNMP request of chassis contents. [Defect ID 44526]
- High and abated memory utilization traps incorrectly swap values. [Defect ID 47967]

SONET

- ATM SONET bit error counters are not incremented for ports 1, 2, and 3 of 4xOC3-ATM. [Defect ID 48454]

System

- The SRP module reset with a panic in task: cliLocal file: ConfigChecker.cc line: 847. [Defect ID 49685]
- CLI display resulting from failed temperature sensor is unclear. [Defect ID 49727]
- On an ERX system with a single SRP module, the **halt** command fails with the following message: [Defect ID 47833/47935]

```
% Unable to halt at this time (standby SRP may be disabled or
not present, powering down the chassis while SRP Flash
disk write access is in progress may corrupt the content of
its Flash card)
```
- SRP module reset: Hardware watchdog after modifying boot string from shell in redundant system. [Defect ID 49181]
- SRP module reset: panic, arg (0x10000),file: icc.cc, line: 306, task: icLoader. [Defect ID 49183]
- After deleting multiple virtual routers, not all of NVS cache was flushed. [Defect ID 47165]
- Lost console and telnet connectivity to three routers. [Defect ID 49641]
- SRP module reset: reset type: processor exception 0x300 (data access), task: cm. [Defect ID 42558]
- Message “ERROR 06/12/2003 15:15:33 fileSystem (): A corrupted file was found on the router” while configuring the router. [Defect ID 48637]
- May get ERROR level log when deleting a release. [Defect ID 48857]
- The **show running** command modifies configuration. [Defect ID 46466]
- Files storing halt reset information will not be saved on redundant SRP system. [Defect ID 47596]
- Software support required for new version of the OC48 I/O module. [Defect ID 47145]

System Maximums



Table A-1 presents current system maximums for various hardware configurations. The system does not simultaneously support all maximum configurations.

Modules referred to in the table are identified by their physical label. See *ERX Installation and User Guide, Appendix B, Module Specifications* for a list of modules and their identifying labels.

Table A-1 System maximums for the ERX edge router

Feature	ERX-700 series	ERX-1400 series
Fabric size	10 Gbps	10 or 40 Gbps
Chassis per 7-foot rack	6	3
Physical density providing wire speed performance		
Channelized E1 ports per chassis (CE1 Full I/O modules)	100	240
Channelized OC3 ports per chassis (cOC3 STM1 FO I/O modules); wire rate/shared	16/20	32/48
Channelized OC12 ports per chassis (cOC12 STM4 FO I/O modules); wire rate/shared	4/5	4/12
Channelized T3 ports per chassis (CT3/T3 I/O modules)	15	36
Channelized T3 ports per chassis (CT3/T3 12 I/O modules); wire rate/shared	48/60	96/144
Channelized T1 ports per chassis (CT1 Full I/O modules)	120	288
E3 (unchannelized) ports per chassis (CT3/T3 I/O modules)	15	36
E3 (unchannelized) ports per chassis (E3 12 FRAME I/O modules)	48/60	96/144
Fast Ethernet (10/100) ports per chassis (FE-8 I/O modules); wire rate/shared	32/40	32/96

Table A-1 System maximums for the ERX edge router (continued)

Feature	ERX-700 series	ERX-1400 series
Gigabit Ethernet ports per chassis (GE I/O modules); wire rate/shared	4/5	4/12
HSSI ports per chassis (HSSI-3F I/O modules)	15	36
OC3/STM-1 ATM ports per chassis (OC3-4 I/O modules); wire rate/shared	16/20	32/48
OC3/STM-1 POS ports per chassis (OC3-4 I/O modules); wire rate/shared	16/20	16/48
OC12/STM-4 ATM ports per chassis (OC12 STM4 I/O modules); wire rate/shared	4/5	8/12
OC12/STM-4 POS ports per chassis (OC12 STM4 I/O modules); wire rate/shared	4/5	4/12
OC48/STM16 POS ports per chassis (OC48 FRAME I/O modules); ERX-1440 system only	N/A	2
T3 (unchannelized) ports per chassis (CT3/T3 I/O modules)	15	36
T3 (unchannelized) ports per chassis (4xDS3 ATM I/O modules)	16/20	32/48
T3 (unchannelized) ports per chassis (CT3/T3 12 I/O modules)	48/60	96/144
Logical density per chassis		
Logical E1s per chassis	1,260	3,024
Logical E3s per chassis	60	144
Logical fractional E1s (DS0) per chassis	3,100	7,440
Logical fractional T1s (DS0) per chassis	2,880	6,912
Logical OC3/STM1 per chassis	20	48
Logical OC12/STM4 per chassis	5	12
Logical OC48/STM16 per chassis (ERX-1440 system only)	N/A	2
Logical T1 per chassis	1,680	4,032
Logical T3 per chassis	60	144

Table A-1 System maximums for the ERX edge router (continued)

Feature	ERX-700 series	ERX-1400 series
Logical density per module combination		
Logical E1s per CE1 line module and supported I/O modules	20	20
Logical E1s per cOCx/STMx F0 line module and supported I/O modules	252 63 per OC3/STM1	252 63 per OC3/STM1
Logical E3s per E3 ATM or E3 FRAME line module and supported I/O modules	3	3
Logical E3s per COCX-F3 line module and supported I/O modules	12	12
Logical fractional DS0s per cOCx/STMx F0 line module and supported I/O modules	2,000 500 per OC3/STM1	2,000 500 per OC3/STM1
Logical fractional E1s (DS0) per CE1 line module and supported I/O modules	620	620
Logical fractional T1s (DS0) per CT1 line module and supported I/O modules	576	576
Logical fractional T1s (DS0) per CT3 line module and supported I/O modules	384	384
Logical fractional T1s (DS0) per CT3/T3-F0 line module and supported I/O modules	1992 166 per T3	1992 166 per T3
Logical fractional T3s (DS3) per COCX-F3 line module and supported I/O modules	12	12
Logical fractional T3s (DS3) per T3 FRAME line module and supported I/O modules	3	3
Logical T1s per cOCx/STMx F0 line module and supported I/O modules	336 84 per OC3/STM1	336 84 per OC3/STM1
Logical T1s per CT3 line module and supported I/O modules	84 28 per T3	84 28 per T3
Logical T1s per CT3/T3-F0 line module and supported I/O modules	336 28 per T3	336 28 per T3
Logical T3s per COCX-F3 line module and supported I/O modules	12	12
Logical T3s per cOCx/STMx F0 line module and supported I/O modules	12 3 per OC3/STM1	12 3 per OC3/STM1
Logical T3s per CT3/T3-F0 and supported I/O modules	12	12
Logical T3s per OCx/STMx ATM line module and 4xDS3 ATM I/O module	4	4
Logical T3s per T3 ATM or T3 FRAME line module and supported I/O modules	3	3

Table A-1 System maximums for the ERX edge router (continued)

Feature	ERX-700 series	ERX-1400 series
ATM virtual circuits per chassis	64,000	64,000
ATM virtual circuits per line module		
E3 ATM	4,000	4,000
OC3	4,000	4,000
OCx/STMx ATM (active/configured)	8,000/16,000	8,000/16,000
T3 ATM	4,000	4,000
ATM virtual circuits per port		
E3 ATM	4,000	4,000
OC3	4,000	4,000
OCx/STMx ATM (active/configured)	8,000/16,000	8,000/16,000
T3 ATM	4,000	4,000
ATM VP/VC addresses per line module		
E3 ATM	18-bit	18-bit
OCx/STMx ATM with an OC3-4 I/O module	20-bit	20-bit
OCx/STMx ATM with an OC12 STM4 I/O module	20-bit	20-bit
OCx/STMx ATM with a 4xDS3 ATM I/O module	20-bit	20-bit
T3 ATM	18-bit	18-bit
ATM VP tunnels per port, all modules	256	256
Dynamic Interfaces		
Active autosensed dynamic interface columns per chassis	32,000	32,000

Table A-1 System maximums for the ERX edge router (continued)

Feature	ERX-700 series	ERX-1400 series
Ethernet density		
VLANs (FE-2 I/O)	4,096	4,096
VLANs (FE-8 I/O)	8,192	8,192
VLANs (GE I/O)	4,096	4,096
VLANs per chassis	32,768	32,768
S-VLANs (FE-2 I/O)	4,096	4,096
S-VLANs (FE-8 I/O) (spread across at least two ports)	8,192	8,192
S-VLANs (GE I/O)	8,192	8,192
S-VLANs per chassis	32,768	32,768
Frame Relay virtual circuits per chassis		
	5,000	12,000
Frame Relay virtual circuits per port		
CE1	1,000	1,000
COCX-F3	1,000	1,000
cOCx/STMx F0	1,000	1,000
CT1	1,000	1,000
CT3	1,000	1,000
E3 FRAME	1,000	1,000
T3 FRAME	1,000	1,000
MLPPP bundles per chassis		
	4,000	4,000
PPP sessions per chassis		
	32,000	32,000
Frame Relay virtual circuits per line module		
CE1	1,000	1,000
COCX-F3	1,000	1,000
cOCx/STMx F0	1,000	1,000
CT1	1,000	1,000
CT3	1,000	1,000
E3 FRAME	1,000	1,000
T3 FRAME	1,000	1,000

Table A-1 System maximums for the ERX edge router (continued)

Feature	ERX-700 series	ERX-1400 series
PPP sessions per line module		
CE1	620	620
COCX-F3	12	12
cOCx/STMx FO	2,000	2,000
CT1	620	620
CT3	384	384
E3 ATM	4,000	4,000
E3 FRAME	4,000	4,000
FE-2	4,000	4,000
GE/FE	8,000	8,000
OC3	4,000	4,000
OCx/STMx	8,000	8,000
OC48/STM16 (ERX-1440 system only)	N/A	1
T3 ATM	4,000	4,000
T3 FRAME	4,000	4,000
Forwarding Table Entries^a		
Chassis with only ASIC modules	1,048,576	1,048,576
Chassis with one or more non-ASIC modules	65,536	65,536
IP network interfaces		
Per chassis	32,000	32,000
Per line module ASIC/non-ASIC	8,000/4,000	8,000/4,000

Table A-1 System maximums for the ERX edge router (continued)

Feature	ERX-700 series	ERX-1400 series
Routing protocol scaling and peering densities^b		
Routing table entries	500,000	500,000
BGP-4 peering sessions	1,000	1,000
BGP-4 routes (NLRI)	1,500,000	1,500,000
IP/MPLS next hops (egress FECs) on system with at least one non-ASIC module	256,000	256,000
IP/MPLS next hops (egress FECs) on system with ASIC modules only	1,000,000	1,000,000
IS-IS adjacencies	150	150
IS-IS routes	10,000	10,000
Layer 2 circuits over MPLS per line module	8,000 ^c	8,000 ^c
Layer 2 circuits over MPLS per chassis	10,000	10,000
MPLS CR-LDP LSPs	8,000	8,000
MPLS LDP LSPs	8,000	8,000
MPLS RSVP-TE LSPs	4,000	4,000
OSPF adjacencies	150	150
OSPF routes	10,000	10,000
Performance		
IP packet forwarding (packets per second)	7,500,000	18,000,000
Policy and QoS		
QoS queues per ASIC line module	48,000	48,000
Classification rules per policy	512	512
Policy classification (CLACL) entries per line module (ASIC and non-ASIC)	256,000	256,000
Unique policy assignments per line module (ASIC and non-ASIC)	8,159	8,159
Policy egress interface attachments per ASIC line module	8,191	8,191
Policy ingress interface attachments per ASIC line module	8,191	8,191
Policy interface attachments per non-ASIC line module	8,191	8,191
Rate limiters (egress) per ASIC line module	24,575	24,575
Rate limiters (ingress) per ASIC line module	24,575	24,575
Rate limiters per non-ASIC line module	24,575	24,575
Policies statistics blocks (egress) per ASIC line module	65,535	65,535
Policies statistics blocks (ingress) per ASIC line module	65,535	65,535
Policies statistics blocks per non-ASIC line module	32,767	32,767
Software lookup blocks per line module (ASIC and non-ASIC)	1,023	1,023

Table A-1 System maximums for the ERX edge router (continued)

Feature	ERX-700 series	ERX-1400 series
VPN/ tunnels		
DVMRP (IP in IP) tunnels per chassis	4,000	4,000
DVMRP (IP in IP) tunnels per Tunnel Service Module	4,000	4,000
GRE tunnels per chassis	4,000	4,000
GRE tunnels per Tunnel Service Module	4,000	4,000
IPsec tunnels per IPSec service line module	5,000	5,000
IPSec manual secure tunnels per chassis	256	256
IPSec transform sets per chassis	1,000	1,000
IPSec transforms per transform set	6	6
IPSec tunnels per chassis	10,000	10,000
L2F tunnels per module and per chassis	200	200
L2F sessions per module and per chassis	2,000	2,000
L2TP LAC sessions per chassis	16,000	16,000
L2TP LAC tunnels per chassis	4,000	4,000
L2TP LNS sessions per chassis	16,000	16,000
L2TP LNS tunnels per chassis	4,000	4,000
L2TP LNS sessions per Tunnel Service Module	8,000	8,000
L2TP LNS tunnels per Tunnel Service Module	4,000	4,000
Virtual routers per chassis	1,000	1,000
Virtual routers per line module ASIC/non-ASIC	1,000/250	1,000/250

- a. The total set of FTEs can be shared by interfaces, next hops, ECMP sets, VRs, and VRFs. Next-hop FTEs identify the next hop on multiaccess media, such as ATM multipoint, Ethernet, or bridged Ethernet. Each VR or VRF consumes 3 entries. Each interface, next hop, and ECMP set consumes a single entry. One FTE is reserved for internal use and the system software limits the number of FTEs used by interfaces to a maximum of 32,000. The remaining FTEs can be shared across the other types.
- b. These values are subject to limitations on available SRP module memory, which varies according to your system configuration.
- c. This maximum is not valid for Frame Relay. The Frame Relay maximum is 1,000 circuits over MPLS per line module, because only 1,000 Frame Relay DLCIs are permitted per line module.