



ERX Edge Routers Release Notes

Release 4.1.3

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

Part No. 162-00486-05 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. ERX, ESP, E-series, Internet Processor, J-Protect, JUNOS, JUNOScript, JUNOSe, M5, M7i, M10, M10i, M20, M40, M40e, M160, M-series, NMC-RX, SDX, T320, T640, and T-series 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 ERX-310, ERX-705, ERX-710, ERX-1410, ERX-1440, M5, M7i, M10, M10i, M20, M40, M40e, M160, and T320 routers, T640 routing node, and the JUNOS, JUNOSe, and SDX-300 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, 6,192,051, 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.

ERX Edge Routers Release Notes, Release 4.1.3
Copyright © 2004, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writers: Brian Wesley Simmons, Fran Singer
Editor: Fran Mues

Revision History
March 2004

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 NONINFRINGEMENT 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 NONINFRINGEMENT.

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	5
Contacting Customer Service	6
Release Highlights	6
Release Software Protocols	7
Core Routing Stack	7
Routing Protocols	7
Multiprotocol Label Switching (MPLS)	8
Layer 2 Protocols	8
Security Protocols	8
Simple Network Management Protocol (SNMP)	8
Known System Behavior	9
ATM	9
BGP	9
B-RAS	10
CLI	10
Ethernet	10
Hardware	10
IP	11
L2TP	11
Layer 2 over MPLS	12
MLPPP	12
MPLS	12
Policy Management	13
QoS	13
Syslog	13
System	14

Known Problems and Limitations	14
ATM	14
B-RAS	15
CLI	15
Documentation	15
GRE	16
IP	17
MLPPP	18
MPLS	18
NTP	18
OSPF	19
QoS	19
SDX	19
System	19
Resolved Known Problems	20
AAA	20
ATM	20
BGP	21
BGP/MPLS VPNs	22
Bridged Ethernet	22
CLI	22
Core Dump	22
DHCP Local Server	22
DHCP Relay	23
Diagnostics	23
Ethernet	23
Firmware	23
Forwarding	23
Frame Relay	25
Hardware	25
HDLC	25
IP	25
IP Multicasting	26
IS-IS	27
L2TP	27
MPLS	27
NTP	29
OSPF	29
PIM-SM	29
Policy Management	30
PPP	30
PPPoE	30
QoS	30
RADIUS	30
Redundancy	30
RIP	31
Routing Policy	31
SDX	31
SNMP	32
SSH	32

System	32
T3	34

Appendix A System Maximums

ERX 4.1.3 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 4.1.3 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	9
Known Problems and Limitations	14
Resolved Known Problems	20
<i>Appendix A, System Maximums</i>	A-1

Before You Start

These *Release Notes* include information on the changes between Releases 4.1.2 and 4.1.3. 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 4.1.3 Release Notes, which describe features available in Release 4.1.3.
- The 4.1.0, 4.1.1, and 4.1.2 Release Notes, which describe features available in Releases 4.1.0, 4.1.1, and 4.1.2.
- 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 4.1.3 is a maintenance release, but does include the following hardware feature.

- Support for 1-GB SRP-5G+ and SRP-10G Modules

New 1-GB versions of the following SRP modules are now supported:

- > 1-GB SRP-5G+ module for installation in an ERX-705 router
- > 1-GB SRP-10G module for installation in an ERX-710 or ERX-1410 router

You can identify the new SRP modules as follows:

Model Number	Description	Assembly Number
ERX-5G1GECC-SRP	1-GB SRP-5G+ module	350-00053-51 or later
ERX-10G1GECC-SRP	1-GB SRP-10G module	350-00054-52 or later

Each of these SRP modules has been upgraded with the following memory and storage capabilities to support specific software features available in future JUNOSe releases:

- > 1-GB of memory with error checking and correction (ECC)
- > 1-GB Type II PCMCIA nonvolatile storage (NVS) card

In addition, the 1-GB SRP-5G+ module and the 1-GB SRP-10G module each pairs with an enhanced SRP I/O module (Assembly Number 350-00033-02 Rev. A01 or later) that is required to support certain software features in future JUNOSe releases.

You must install one of the following minimum JUNOSe software releases to use the 1-GB SRP-5G+ or 1-GB SRP-10G module:

- > Release 4.1.3 or later 4.1.x release
- > Release 5.0.4 or later 5.0.x release

- > Release 5.1.2 or later

Once you install the 1-GB SRP-5G+ or 1-GB SRP-10G module, the JUNOSe software prevents you from installing a release that is earlier than these minimum required releases.

For complete instructions on installing the 1-GB SRP-5G+ or 1-GB SRP-10G module in an E-series router, see the notice *Upgrading Memory on SRP-5G+ and SRP-10G Modules* (Part Number 162-00835-00 Rev. A00), which comes with the new module.

- Support for Mixed SRP Modules

JUNOSe supports a physical configuration that includes both a 512-MB and a 1-GB SRP module in a redundant configuration with the following restrictions:

- > The system configuration must be compatible with a 512-MB SRP module.
- > The 512-MB SRP module must be configured as primary and the 1-GB SRP module as redundant.
- > If SRP failover makes the 1-GB SRP module primary, problems may occur if the configuration changes while the router is in this state.

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.

Syslog

- If you enable engineering logs and set the control network logs to a level of “notice” or lower (down from the default of “error”), you might see erroneous controlNetwork log messages like the following that are

generated due to SNMP polling on line modules (correctly) detecting no fabric: [Defect ID 43168]

```
NOTICE 09/01/2002 18:47:52 CEST controlNetwork (slot 11):  
Control Bus Master slave error 0x5 while accessing slot
```

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.

- In some circumstances, the router behaves as though configuration changes have been made that need to be committed to NVS, when no such changes have been made. This behavior occurs, for example, if you issue the following sequence of commands: **service manual-commit**, **write memory**, **show running-configuration** and subsequently issue the **reload**, **srp switch**, **halt**, or **show environment** command. [Defect ID 48058]

Work-around: Issue the **write memory** command when the CLI indicates that it is necessary, even if you believe no changes have been made.

- 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, , 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 OCx/STMx ATM and OCx/STMx 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.

GRE

- When you delete a virtual router that has been configured as a transport virtual router for either a GRE or DVMRP tunnel, the show configuration output displays No Router for the transport virtual router. [Defect ID 44810]

Work-around: To remove such a tunnel interface, simply omit any reference to the transport virtual router. For example, to delete

interface tunnel gre:v6Tunnel transport-virtual-router No Router from the configuration, issue the command, **no interface tunnel gre:v6Tunnel**.

IP

- If you issue the **ip tunnel reassembly** command on a virtual router, fragmented L2TP control packets are discarded. The `ddShowVrReasmStats` output on the TSM shell reports and “unrecognizedProtocolError” for the fragmented control packets. The TSM reassembles L2TP data packets without any trouble. [Defect ID 54932]
- 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. [Defect ID 3012]

Work-around: Display the local and peer endpoint discriminators via the **show mlppp interface full** command to help debug the problem.

MPLS

- If LSPs are announced into IGP, then the IGP routes cannot be used for multicast RPF checks, because LSPs are unidirectional. [Defect ID 28526]

Work-around: Configure static RPF routes with native hops if LSPs are auto-route announced to IGP.

NTP

- The source address for outgoing NTP packets is not updated with the configured address. [Defect ID 45619]
- If you configure NTP with a master stratum number, the configuration is not displayed by the show configuration command. [Defect ID 47012]
- If you issue the **no ntp enable** command after one or more NTP server associations have already been configured, the association configuration disappears from the output of the **show configuration** command. [Defect ID 47036]
- 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]

SDX

- The SDX client intermittently reports an incorrect interface count via the **show ssc info** command and SNMP. [Defect ID 51544/52356]

System

- During line module failover and reversion testing, the SNMP description that was “OC3 quad port, ATM” before failure becomes OCX ATM. [Defect ID 51364]

Work-around: Reload the router or switch the SRP modules to reinstate the correct SNMP description.

- 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. [Defect ID 34842]
- 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.
- 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 4.1.2p1-10 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

- Start-acct messages aren't being sent when using a non-PPP equal access. [Defect ID 50443]
- Router generates Acct-Stop on Access-Reject. [Defect ID 52711]
- The **show subscribers** command uses the wrong “interface” value for non-tunneled users. [Defect ID 51732]

ATM

- Redirecting output of **show atm vc** command to a file results in a wrong VC count in the display. [Defect ID 55411]
- After issuing the **erase secrets** command, you cannot execute the **enable** command if **aaa authentication enable default** is set. [Defect 54940]

- Dynamic interfaces not fully recovering on various port/line module flaps. [Defect ID 55627]
- Line module reset: exception 0x68616c74 (halt), task: scheduler. [Defect ID 52124]
- The **atm vc-per-vp** command incorrectly denied “% Cannot set vc-per-vp range (circuit exists)”. [Defect ID 52087]
- L2TP ICCN Connect Speed AVP shows speed in Kbps rather than bps as per spec. [Defect ID 51729]
- Egress forwarding on ATM multipoint interface stops after deleting and reconfiguring VCs. [Defect ID 50194]
- ATM Vc-per-Vp configuration is not applied correctly after a line module failure. [Defect ID 51585]
- After an upgrade from 4-0-3 to 4-1-1p1-3 the **atm oam flush alarm-cell** command has changed to **atm oam flush**. [Defect ID 50620]
- Line module reset: type: panic file: 83DataService.cc line: 5606 task: scheduler. [Defect ID 51417]
- Router reset: processor exception 0x300 (data access), task: telC7712. [Defect ID 50342]
- NBMA does not seem to be working properly: only 2 out of 3 destinations receive a copy of packet. [Defect ID 50642]

BGP

- BGP removes VPNV4 routes even when automatic route target filtering is disabled. [Defect ID 54803]
- SRP module reset type: panic task: bgp file: bgpNextHop.cc line: 5714 when adding and deleting VRFs. [Defect ID 50366]
- BGP VPNs: should gracefully handle a received explicit null label -value 0 and 2 for VPN routes. [Defect ID 51449]
- Reset in BGP processor exception 300 when deleting a IP VRF after deleting a large number of import Route Targets. [Defect ID 50806]
- BGP peer OID 1.3.6.1.2.1.15.3.1.14 (bgpPeerLastError) returning incorrect value. [Defect ID 50769]
- SRP module reset at bgpNexthop.cc at line 5261 task: bgp2. [Defect ID 50653]
- External type 2 is preferred over external type 1. [Defect ID 50631]

BGP/MPLS VPNs

- BGP/MPLS VPN import map filters do not work for routes leaked in from other VRF's on the local PE router. [Defect ID 53230]
- AS-path information is incorrect using **neighbor as-override** command in a BGP confederation. [Defect ID 53986]

Bridged Ethernet

- On the ERX-1440 and ATM module reset due to processor exception 0x68616c74 (halt); task: scheduler; when the **shut** command was issued. [Defect 55050]
- SRP module reset with processor exception 0x300 (data access) task: bridgedEthernet. [Defect ID 53562]

CLI

- Percent running time process statistics (**show process cpu** command) are incorrectly calculated. [Defect ID 27865]
- Tabular report with rows that wrap and report is longer than 1 terminal page (**show process** command) may miscalculate column widths. [Defect ID 50147]
- Cannot remove files of type .scr from /incoming and /outgoing directories. [Defect ID 44240]
- Arrow key on command line is not behaving correctly. [Defect ID 51590]
- The **show configuration** command output much slower with 4-1-1p1-3 than with 4-0-2p0-5. [Defect ID 51510]

Core Dump

- 2GB SRP module core dump takes 2 hours to complete. [Defect ID 49864]

DHCP Local Server

- DHCP-LS does not send an OFFER upon receiving a DISCOVER in the stateSelectingWithAddress state. [Defect ID 50914]
- ERX DHCP server ignores DHCP request if a DHCP ACK is missed. [Defect ID 54564]

DHCP Relay

- SRP module reset with exception at task:system. [Defect ID 53027]
- DHCP relay statistics are not properly incrementing when receiving packets with the giaddr set to the router's address. [Defect ID 52302]

Diagnostics

- Remove GE cam tests from POST and test CLI. [Defect 51919]
- Spurious line module core dump created after a normal reset. [Defect ID 52293]
- TSM module stuck in booting after diagnostic failure. [Defect ID 51392]
- Received NP55 errors messages when standby SRP failed to sync. [Defect ID 51885]
- Change “read-verify” command to “read” in BPOST and Scandisk. [Defect ID 46290]

Ethernet

- Reg (FEPHY): FastEthernet connections do not recover after flapping interface. [Defect ID 52535]
- Router LNS subscriber login problems. [Defect ID 49860]
- DEBUG code for: Line module reset: DPFE modules reset intermittently due to Ping Failure Threshold Exceeded. [Defect ID 51622]

Firmware

- POS interface not receiving PPP control traffic. Reload of line module restores connection. [Defect ID 51370]
- Line module reset: ping failure threshold exceeded resets. [Defect ID 50736]

Forwarding

- Line module reset with UT3A unknown software error signature (0xead0000) in task: scheduler, file ut3aSar.cc, line 794. [Defect ID 54037]
- PPP users can not log into ERX, ERX configured as an LNS. [Defect ID 51322]

- FE-8 line module resets with exception 0x100 after shutting interface in IP integration network. [Defect ID 51955]
- Martini mpls frame-relay compatibility issues when connected to a Cisco. [Defect ID 53802]
- The **show ip route slot** command displays only a subset of the routes even when max paths set to 16 on non-ASIC modules. [Defect ID 54166]
- Panic in geFcDriver.cc line 802 when shutting interface; subsequent halt in IP. [Defect ID 49951]
- Counters on VLAN interfaces are not accurately reflecting the traffic passing over them. [Defect ID 51725]
- DPFE reset type: processor exception 0x300 (data access) task: scheduler -> fc1EthernetIngressProcessing. [Defect ID 52370]
- FC Recovery type reset incorrectly recorded for 1st generation line modules. [Defect ID 50925]
- Reassembly failures on TSM in L2TP configuration, router as LNS. [Defect ID 51533]
- Line module reset: ping failure threshold exceeded reset. [Defect ID 50687]
- An interface configured with a Rate Limiting Policy stops forwarding traffic at random intervals. [Defect ID 49678]
- Spurious, invalid line module core dump produced after a “normal” reset. [Defect ID 50664]
- FC reset initiated when ARF returns out-of-range value. [Defect ID 51746]
- Line module reset: GE, FC processor exception 0x700 (program), task: Egress FC[1], fpEgressProcessWhiteRouteTag Mok_ERX_2. [Defect ID 49400]
- Problems forwarding traffic on a 4 port OC3 line module. [Defect ID 50142]
- OC48 ECO: change base clock frequency to 50MHz. [Defect ID 49029]
- Line module reset: Ping failure threshold exceeded reset while running traffic over subscriber interfaces. [Defect ID 49612]
- 10G fabric discards cells. [Defect ID 50577]
- ATM throughput performance reduction in 3-4-1. [Defect ID 50353]

Frame Relay

- UT3F-12-reset type: panic task: scheduler file: frameRelayLmi.cc line: 897. [Defect ID 49038]

Hardware

- Line module reset with processor exception 0x68616c74 (halt), task: scheduler. [Defect ID 50246]
- Line module reset: Processor exception 0x300 (data access), task: scheduler. [Defect ID 52070]
- FC recovery could cause a panic in ifaClassifier.cc or ifaIngress.cc. [Defect ID 53298]

HDLC

- Line module reset with processor exception 0x68616c74 (halt) task: scheduler halter PC: 0x275278 pc: 0x33c520 -> EfaSchedulerCore::drainPort(int). [Defect ID 35880]
- 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]

IP

- The **show ip forward slot x** command indicates updating state for virtual routers that do not have interfaces on slot x. [Defect 55407]
- Mroutes with null OIF list not installed on the line moduole. [Defect 51182]
- Forwarding tables stuck in updating mode. [Defect 55212]
- RIP routes with a distance of 255 are installed in the IP routing table. [Defect ID 53803]
- IP routing loop in IS-IS/LDP topology-driven environment. [Defect ID 52269]
- Proxy-ARP does not work as expected. [Defect ID 52264]
- Secondary IP addresses on router interfaces inherit the layer 2 MTU instead of the administratively configured IP MTU. [Defect ID 52871]

- The **no ip summary-address** command takes long time to appear in in **show configuration** output. [Defect ID 53590]
- ICMP redirect sent instead of frag needed but DF set. [Defect ID 53892]
- SRP module reset with exception task telC3 while add/del 1023 VRs with multiple sub-interfaces across different LCs. [Defect ID 50521]
- Receiving log messages: ipProfileMgrEngineering ():
AaaInfo::setLocalInterface: address/mask format invalid: Loopback 2. [Defect ID 51479]
- Deleting a static route via SNMP appears to work but does not delete the entry in NVS. [Defect ID 51737]
- SRP module reset with processor exception 0x68616c74 (halt), task: ip_RxData_1 deadlock. [Defect ID 50361]
- RPF check returns LSP causing multicast stream to be discarded. [Defect ID 51165]
- SRP module reset with panic task: ip_Ctrl_1AD file: ar1IpEngine.cc line: 7727 When adding and deleting VRFs. [Defect ID 50369]
- The **show ip route slot** command does not display ECMP routes. [Defect ID 50815]
- The **show ip route slot** command fails although the route does exist on the line module. [Defect ID 51112]
- SRP module reset with exception 0x68616c74 (halt) while upgrading from 4-1-1p1-2 to 4-1-2. [Defect ID 50883]
- If there are multiple loopback interfaces configured, the **show configuration** and **show ip interface** commands do not work. [Defect ID 51211]
- SRP module reset with 0x300 data access exception in task: ip_Ctrl_245 after issuing **ip si0** command. [Defect ID 52709]
- The **show ip route slot x** command causes line module to reset due to panic; file: IngressFecApi.cc; line: 71; task: scheduler. [Defect ID 54545]

IP Multicasting

- JUNOSe does not enforce DA=GA for IGMP membership reports (“joins”). [Defect ID 55133]
- Line module reset in ipTunnel.cc on line 1391 in task scheduler. [Defect ID 49442]

- IGMP limit. [Defect ID 51698]
- The **log severity debug igmpGeneral** command does not log any IGMP information. [Defect ID 50431]
- Initial multicast packet for a multicast stream may be dropped by the egress line module. [Defect ID 50812]
- Exception 0x300 (data access) task: mgtm_ctrl_4. [Defect ID 50201]

IS-IS

- IS-IS metric value incorrect when using wide metric. [Defect ID 54456]
- SRP module reset with processor exception 0x100 (watch dog), task: icc. [Defect ID 53386]
- SRP reset with panic, arg (0x28) file: ipTypes.cc line: 143 task: isis1r1. [Defect ID 51321]
- IS-IS does not correctly install all next hops of an ECMP set. [Defect ID 51010]

L2TP

- JUNOSe LNS implementation returns invalid session ID in CDN control packet. [Defect ID 54842]
- SRP module reset with panic, task: l2tp, file: icatorSession.cc, line: 1417. [Defect ID 53780/51647]

MPLS

- SRP module reset with panic, task: mplsMgr1r2, file: rsLdpFecInfo.cc, line: 1564. [Defect ID 55732]
- Failure to create mplsMinorInterface in one VRF due to lack of router ID in others. [Defect ID 55640]
- 5-1-0b2-18;Reg (ipReg); ERROR;Transit to Egress Merge Failed Delete Fec Uid: MplsMinor 0x1a000f20 Already an Egress. [Defect ID 51859]
- SRP module resets with a panic in task: mplsMgr3r129, file: osPool.cc, line: 1426. [Defect ID 54634]
- The **show configuration** command is hanging as well as certain IP related commands due to deadlock. [Defect ID 52491]

- If you explicitly configure MPLS on an interface that is an inter-AS BGP/MPLS VPN interface, the inter-AS VPN will not work. [Defect ID 47067]
- Interoperability issues between GSR/E-series router while running MPLS. [Defect ID 53052]
- LDP convergence is slow in JUNOSe 4-1-2. [Defect ID 52000]
- LDP loop detection is incorrectly enabled by default in JUNOSe for topology driven LSPs. [Defect ID 51999]
- MPLS LSPs not built over VLANs after a topology change. [Defect ID 50886]
- Multi-telnet/add remove VRF pairs causes exception 0x300 in mplsMgr450r901 (TopologyDrivenLdpEnabled). [Defect ID 50578]
- Transit LSP not built after OSPF is changed to use ECMP. [Defect ID 51060]
- vir rou 'CE1' reset type: exception 0x400 (instruction access); task: mplsMgr24r134. [Defect ID 50513]
- Routes are not populated in the VRF table (for a set of VRFs). [Defect ID 50576]
- VPN over LDP on UT3F/PPP interfaces sometimes does not recover after line module reload. [Defect ID 46276]
- SRP module reset with processor exception 0x300 (data access) task: mplsMgr1r1. [Defect ID 50990]
- MPLS does not converge on a new IS-IS path when the IS-IS metrics of the two paths are not equal. [Defect ID 51038]
- 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]

NTP

- NTP does not synchronize after a reload. [Defect ID 53231]

OSPF

- OSPF external routes preferring a stub area path over a backbone area path. [Defect ID 55861]
- SRP module reset with panic in file: osPool.cc line: 1689 task: ospf_msgQueue_3. [Defect ID 55666]
- OSPF is not selecting the correct external type-1 route after an SPF calculation. [Defect ID 54966]
- Rewrite of metric type in route map not enforced with summary external routes in OSPF. [Defect ID 52554]
- SPF tree calculation incorrectly considers link entries in a router LSAs with a cost of 65535 as unreachable. [Defect ID 52933]
- OSPF route selection for external type 1 routes is incorrect. [Defect ID 50810]
- The **show ip ospf database external** command does not display all external routes associated with specific link state-id. [Defect ID 51094]
- Router reset with panic in ipNextHopInfo.cc line 65. [Defect ID 50093]
- OSPF adjacencies stuck in loading after a duplicate router ID is removed from the network. [Defect ID 51645]

PIM-SM

- PIM join from Nortel ARN is being rejected due to nonzero data in a reserved field. [Defect ID 51619]
- SRP module reset with processor exception 0x300 (data access), task: pim1 when it learns 10,000 mcast groups. [Defect ID 52144]
- Multicast (PIM-SM) ERX doesn't switch from a shared tree to the SPT after the S,G has been joined, left, and joined again. [Defect ID 50151]
- PIM-SM auto-rp mapping agent not sending selected RP information to other routers. [Defect ID 51942]

Policy Management

- Policy management erroneously tests for a maximum policy rule number. [Defect ID 44875]
- Unable to mark ToS bits with rate-limit-profile. [Defect ID 52620]
- Baseline IP interface statistics doesn't work for output policy. [Defect ID 52679]
- CT3 module reset with panic; file: cPolicyDriver.cc; line: 2577; task: scheduler. [Defect ID 51927]

PPP

- Static PPPoE subinterface over bridge1483 interface appears in all virtual routers. [Defect ID 53702]
- Line module reset with panic, msg “6500 10000235 c0210915”, file: ppp.cc, line: 5623. [Defect ID 53242]

PPPoE

- FE-8 module reset with (ic1Fc2L2Api.cc, line 1000) on LAC while bringing PPPoE/L2TP client connections up and down. [Defect ID 47284]

QoS

- QoS queues increment in 512 bytes which may result in traffic latency. [Defect ID 50879]
- Telnet session hangs when **clear egress-queue tunnel mpls:xxxxx** command is issued. [Defect ID 50484]

RADIUS

- The **no virtual-router vrouter1** command fails to delete the VR when a RADIUS acct server is still configured in the VR. [Defect 52394]
- E-series router sending Accounting-Stop packet following SRP switch. [Defect ID 51781]

Redundancy

- System reset: panic in task system; file: lRedundancyDb.cc on line 814. [Defect ID 46583]

- Adjacent redundancy groups do not work in shadow chassis. [Defect ID 50075]

RIP

- RIP not installing default route learned from RIP due to neighbor configuration. [Defect ID 55836]
- SRP module reset with exception 0x300 (data access); task: rip_MsgQ_8000000. [Defect ID 52372]
- RIP neighbor command sends initial advertisement in version 1, even if config tells it should use version 2. [Defect ID 52549]
- SRP module reset with panic, file: ripSummary.cc, line: 151, task: rip_ripMsgQueue_ due to rip summary-address. [Defect ID 53235]
- Interface address not advertised in RIP when IP address is changed. [Defect ID 53140]
- RIP routing table for summary-address disappears by 'clear ip route *'. [Defect ID 53591]
- Redistribution of routes does not work under some situations. [Defect ID 53764]
- SRP module reset with exception 0x300 (data access) task: rip_MsgQ_8000000. [Defect ID 51558]

Routing Policy

- When “set metric +50” is applied to a BGP route without a MED, the result should be a route with MED 50 (now it is “no MED”). [Defect ID 53170]
- SRP module stuck in booting following upgrade from 4-0-3p0-7 to 4-1-2p0-5. [Defect ID 51775]

SDX

- SRP module reset with exception 0x300 (data access) task: cops_DataCtrl_1. [Defect ID 54433]
- SRP module reset with processor exception 0x300 (data access), task: sscClient_Data_1. [Defect ID 52645]
- SIT_main_44:Reg: SRP assertion while running MibWalker Test file: sMibProcUtils.cc line: 439. [Defect ID 50087]
- SRP module reset with exception 0x300 (data access), task: cops_DataCtrl_1. [Defect ID 51909]

- SRP module reset with exception in task: sscClient_Data when running config script. [Defect ID 49711]

SNMP

- SNMP wALK of juniPolicyMIB fails with Generic error. [Defect 54854]
- Sets to disable juniCliSecurityTrapEnable and juniFileXferTrapEnabled via SNMP fail. [Defect ID 52333]
- Incorrect default value of pingCtlFrequency with snmp set provisioning. [Defect ID 52555]
- Traps for PIM are enabled under category snmp. [Defect ID 53614]
- Using JuniFileXferTableEntry to download software causes the MIB agent to be blocked until the download is complete. [Defect ID 50390]
- SRP module reset with panic in task snmp_4l_init; file: RouterAgentNv.cc on line 1080. [Defect ID 53855]
- UT3 line modules unable to boot when SNMP is polling SRP module CPU utilization. [Defect ID 53296]
- SRP module reset with panic in task snmp_DataNctrl_4, file RouterAgentV3.cc on line 3647. [Defect ID 54140]
- Panic in snmpRouterAgentNv.cc line 1962 in booting following upgrade from 4-0-3p0-7 to 4-1-2p0-5. [Defect ID 51924]
- Agent snmp_DataNctrl_1 and agent snmp_DataNctrl_2 are in a deadlock state which caused a vty sessions to hang. [Defect ID 51310]
- SRP module reset with processor exception 0x300 (data access), task: agent1, pc: 0x3be5e50 -> bcrBGO. [Defect ID 51463]
- There is a scenario where different ERX's can have the same snmpEngineID. [Defect ID 52214]

SSH

- Unable to disconnect or clear SSH sessions on router and no new SSH sessions allowed. [Defect ID 53087]

System

- 1000 VRF scaling test takes significantly longer to complete with 5.0.3 vs. 5.1.0. [Defect ID 53589]

- SRP module reset with panic in 0x300 (data access) in sonet_DataNctrl_, sonet.cc line 1218. [Defect ID 53716]
- SRP module reset with assertion, task: cliLocal, file: osTimer.cc, line: 194, last errno: 0x3d0004 following the “reload” command. [Defect ID 52317]
- There are uncommitted files after a reload when release/config checking are enabled. [Defect ID 52550]
- VRF cannot be created via SNMP. [Defect ID 52526]
- LC reset: processor exception 0x300 (data access), task: scheduler. [Defect ID 51539]
- Router resets or generates a startup synchronization error when you issue the **synchronize** command while the SRP module is booting the operational image. [Defect ID 53111]
- “Failed to rename file local:... to linkFs:...” error during ReleaseChecking recovery. [Defect ID 53760]
- SRP module reset with processor exception 0x300 (data access) in task: syncer - during a sync operation. [Defect ID 51278]
- Double entry for same reset in reboot history file. [Defect ID 45706]
- Copying a dump file from standby to user space does not work. [Defect ID 52138]
- SRP-IC gets wedged. SRP stops receiving and transmitting data. [Defect ID 50127]
- Disable background flash checker. [Defect ID 50853]
- SRP module reset with processor exception 0x700 (program), task: ar1InternalNetwo, pc: 0x38ac284. [Defect ID 42039]
- Failure of secondary external timing lead (T1:B) causes primary chassis clock state (T1:A) to fluctuate. [Defect ID 51033]
- CE1/CT3 line module stuck in the initialization state. [Defect ID 51032]
- Forwarding between line modules may be interrupted when changing traffic class from strict-priority to best effort or vice versa. [Defect ID 50981]
- SRP module reset with panic, task: bridge__A2000000, file: 1QueueManager.cc, line: 422, arg: 8. [Defect ID 50821]
- Standby SRP module’s reboot.hty clock off by 5 hours. [Defect ID 50322]

- Panic in file lQueueManager.cc task bridge__A2000000. [Defect ID 45284]

T3

- The **scramble** command does not work for UT3f-12. [Defect 54565]

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.