JunosE™ Software for E Series™ Broadband Services Routers

Release Notes

Release
11.1.2

Published: 2011-01-18

JunosE™ Software for E Series™ Broadband Services Routers Release Notes, Release 11.1.2
Revision History
January 2011—FRS JunosE 11.1.2

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

Software License

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions.

Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details.

For complete product documentation, please see the Juniper Networks Web site at www.juniper.net/techpubs.

END USER LICENSE AGREEMENT

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

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

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

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

   a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.

   b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.

   c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.

e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

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

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

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

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

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

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

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

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

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

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

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

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

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

Table of Contents

Release 11.1.2

Release Installation

Complete procedures for installing the system software are available in *JunosE System Basics Configuration Guide, Chapter 3, Installing JunosE Software*.

New software releases are available for download from the Juniper Networks website at http://www.juniper.net/customers/support. You can use the downloaded image bundle to create your own software CDs.

Before upgrading to a new version of software, save your router'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.

Informational Note: When you upgrade the software on a router that has a large number of interfaces configured, the router might appear to be unresponsive for several minutes. This condition is normal; allow the process to continue uninterrupted.

Upgrading to Release 5.3.0 or a Higher-Numbered Release

When you upgrade from a lower-numbered release to Release 5.3.0 or a higher-numbered release, the higher release might not load if you issue the **boot system** command from Boot mode while the lower-numbered software is running on the router or if you insert a flash card running a higher-numbered release into a system running a lower numbered release. However, if you issue the **boot system** command from Global Configuration mode, the new software loads properly.

Upgrading from Release 5.1.1 or Lower-Numbered Releases to Release 6.x.x or Higher-Numbered Releases

Release 5.1.1 or lower-numbered releases support application images only up to 172 MB. Your software upgrades or application images may be available remotely through Telnet or FTP, or may be delivered on a new NVS card. If you upgrade the JunosE Software using a new NVS card, we recommend you perform the upgrade in two stages: first to an intermediate release and then to the higher-numbered release you want to run. This restriction is not applicable if you upgrade your software remotely through Telnet or FTP.

To install larger application images for Release 6.0.0 and higher-numbered releases, you must first install Release 5.1.2 (or a higher-numbered 5.x.x release). This enables the system to support application images greater than 172 MB. For example, if you are upgrading the software using a new NVS card, you cannot go from Release 5.1.1 to Release 7.2.0 without first upgrading to Release 5.1.2.

See the following table for compatibility of releases.

JunosE Release	Highest Release Able to Load	Cannot Load	Maximum Application Image
5.1.1 or lower-numbered release	5.3.5p0-2 or the highest-numbered 5.x.x release	6.x.x or higher-numbered release	172 MB (approximate)
5.1.2 or higher-numbered release	No limitation	Not applicable	234 MB (approximate)
7.2.0 or higher-numbered release	No limitation	Not applicable	256 MB (approximate)

For more detailed information on installing software, and about NVS cards and SRP modules, see the following documents:

- *JunosE System Basics Configuration Guide, Chapter 6, Managing Modules*
- *Upgrading NVS Cards on SRP Modules in ERX Hardware Guide, Chapter 8, Maintaining ERX Routers*
- *Upgrading NVS Cards on SRP Modules in E120 and E320 Hardware Guide, Chapter 8, Maintaining the Router*

Moving Line Modules Between Releases

The Juniper Networks ERX1440 Broadband Services 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 compatible line module from a Juniper Networks ERX705, ERX710, or ERX1410 Broadband Services Router to an ERX1440 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 ERX1440 router, that router is able to recognize the line module.

If you move a compatible line module from an ERX1440 router to an ERX705, ERX710, or ERX1410 router, the module loads properly in the new router regardless of the release.

SRP Module Memory Requirements

For Release 5.3.0 and higher-numbered software releases on ERX14xx models, ERX7xx models, and the Juniper Networks ERX310 Broadband Services Router, see *ERX Module Guide, Table 1, ERX Module Combinations*, for detailed information about memory requirements.

For Release 8.2.0 and higher-numbered software releases on Juniper Networks E120 and E320 Broadband Services Routers, see *E120 and E320 Module Guide, Table 1, Modules and IOAs*, for detailed information about memory requirements.

Hardware and Software Compatibility

For important information about hardware and software, see the document set as follows:

- Combinations of line modules to achieve line rate performance are in *JunosE System Basics Configuration Guide, Chapter 6, Managing Modules*.

- Compatibility of *ERX router modules with software releases is in ERX Module Guide, Table 1, ERX Module Combinations*.
- Layer 2 and layer 3 protocols and applications supported by *ERX router modules are in ERX Module Guide, Appendix A, Module Protocol Support*.
- Compatibility of E120 router and E320 router modules with software releases is in *E120 and E320 Module Guide, Table 1, Modules and IOAs*.
- Layer 2 and layer 3 protocols and applications supported by IOAs on the E120 router and the E320 router are in *E120 and E320 Module Guide, Appendix A, IOA Protocol Support*.

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

Release Overview

These *Release Notes* cover Release 11.1.2 of the system software for the Juniper Networks E Series Broadband Services Routers and contain the following sections:

If the information in these *Release Notes* differs from the information found in the published documentation set, follow these *Release Notes*.

Before You Start

These *Release Notes* include information about the changes between Releases 11.1.1 and 11.1.2. 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 11.1.2:

- These 11.1.2 *Release Notes*, which describe changes between Release 11.1.1 and Release 11.1.2
- The 11.1.1 *Release Notes* which describe features available in 11.1.1
- The 11.1.x documentation set, which provides detailed information about features available in Release 11.1.0

The 11.1.x documentation set consists of several manuals and is available only in electronic format. You can print your own documentation using the PDF and HTML formats available at the Juniper Networks Technical Documentation Web site at www.juniper.net/techpubs. Refer to the following table to help you decide which document to use:

Task	Document
Install the router	*ERX Hardware Guide* *E120 and E320 Hardware Guide*
Learn about modules	*ERX Module Guide* *E120 and E320 Module Guide* *ERX End-of-Life Module Guide*
Get up and running quickly	*E Series Installation Quick Start poster or ERX Quick Start Guide* *E120 and E320 Quick Start Guide*
Configure the router	*JunosE System Basics Configuration Guide*
Configure physical layer interfaces	*JunosE Physical Layer Configuration Guide*

Task	Document
Configure link layer interfaces	*JunosE Link Layer Configuration Guide*
Configure line module redundancy, stateful SRP switchover, unified ISSU, VRRP, and interchassis redundancy (ICR)	*JunosE Service Availability Configuration Guide*
Configure IP, IPv6 and Neighbor Discovery, and interior gateway protocols (RIP, OSPF, and IS-IS)	*JunosE IP, IPv6, and IGP Configuration Guide*
Configure IP routing services, including routing policies, NAT, J-Flow statistics, BFD, IPSec, digital certificates, and IP tunnels	*JunosE IP Services Configuration Guide*
Configure IP multicast routing and IPv6 multicast routing	*JunosE Multicast Routing Configuration Guide*
Configure BGP, MPLS, Layer 2 service, and related applications	*JunosE BGP and MPLS Configuration Guide*
Configure policy management	*JunosE Policy Management Configuration Guide*
Configure quality of service (QoS)	*JunosE Quality of Service Configuration Guide*
Configure remote access	*JunosE Broadband Access Configuration Guide*
Get specific information about commands	*JunosE Command Reference Guide A to M* *JunosE Command Reference Guide N to Z*
Monitor system events	*JunosE System Event Logging Reference Guide*
Look up definitions of terms used in JunosE technical documentation	*JunosE Glossary*

Release Highlights

Release 11.1.2 is a maintenance release and includes the feature described in this section.

Category	Feature
System	• *Support for Conventional Upgrade on ES2 10G LMs* on page 6

System

- Support for Conventional Upgrade on ES2 10G LMs

  JunosE Software now supports field programmable gate array (FPGA) upgrades on REV-02 ES 210G line modules. You must perform the upgrade using a conventional software upgrade process; you cannot use the unified ISSU process to perform the upgrade. In lower-numbered releases, you were not able to do a conventional software upgrade or a unified ISSU process on REV-02 ES2 10G LMs.

  In this release, design enhancements enable you to perform a stateless upgrade (non-unified ISSU method, with router-wide outage for users) from JunosE releases numbered lower than Release 11.1.2 to JunosE Release 11.1.2. You cannot run a unified ISSU procedure from JunosE releases that do not contain the design changes for control and status register (CSR) FPGA functionality if the router chassis contains ES2 10G LMs or REV-02 ES2 10G LMs.

The following design improvements to FPGA have been made in this release:

- FPGA modifications for resolving fabric-related problems related to SFM redundancy on ES2 4G LMs, ES2 10G LMs, ES2 10G ADV LMs, SRP-100, and SRP-320 modules
- CSR FPGA changes on ES2 10G LMs and REV-02 ES2 10G LMs to enable FPGA upgrades

If your network contains routers installed with LMs other than ES2 10G LMs or REV-02 ES2 10G LMs, a preferred suggestion is to transfer subscribers configured on the ES2 10G LMs or REV-02 10G LMs to other LMs available in your environment when you perform unified ISSU. This transfer of subcribers to other LMs avoids disruption of user sessions owing to the limitation that exists with performing a unified ISSU operation on ES2 10 LMs and REV-02 10G LMs.

Change in existing behavior: Existing feature extended as described here. You can now perform a conventional upgrade from releases numbered lower than Release 11.1.2 to Release 11.1.2, if the router chassis contains ES2 10G LMs or REV-02 versions of the ES2 10G LM fabric slices. In lower-numbered releases, you could only download the image from serial Programmable Read-Only Memory (PROM), which was a factory image. You still cannot perform a unified ISSU operation from releases numbered lower than Release 11.1.2 to Release 11.1.2, if the router chassis contains ES2 10G LMs.

Early Field Trial Features

The features described in this section are present in the code but have not yet been fully qualified by Juniper Networks. These features are available only for field test purposes in this release. If you use any of these features before they have been fully qualified, it is your responsibility to ensure that the feature operates correctly in your targeted configuration.

DHCP

- Support for DHCP External Server, DHCP Local Server, DHCP Relay, and DHCP Relay Proxy on POS Access Interfaces

  The following packet over SONET (POS) module combinations on E Series routers now support configuration of the DHCP external server, DHCP local server, DHCP relay, and DHCP relay proxy applications, alone or in combination, when the POS module is the access interface:

  - POS module combinations on the E120 router and the E320 router:
    - ES2 4G LM with ES2-S1 OC12-2 STM4 POS IOA
    - ES2 4G LM with ES2-S1 OC48 STM16 POS IOA
  - POS module combinations on ERX14xx models, ERX7xx models, and the ERX310 router:
    - OCx/STMx POS line module with OC3-4 I/O module
    - OCx/STMx POS line modules with OC12/STM4 I/O module
    - OC48 line module with OC48 FRAME APS I/O module

  In the current release, this feature is available for early field test purposes only.

You can configure DHCP external server, DHCP local server, DHCP relay, and DHCP relay proxy on these POS modules in either a virtual router (VR) or a VPN routing and forwarding instance (VRF).

As part of this feature, the **pos** keyword has been added to the existing **ip dhcp-local limit** command. To specify the maximum number of IP addresses that the DHCP local server application can supply to all POS access interfaces or to a specific POS access interface, in the range 0–96000, use the **ip dhcp-local limit** command with the new **pos** keyword. For example:

```
! Set the IP address limit for all POS access interfaces to 1000
host1(config)#ip dhcp-local limit pos 1000
! Set the IP address limit for the specified POS access interface to 2000
host1(config)#ip dhcp-local limit interface pos 5/0/0 2000
! Restore the IP address limit for all POS access interfaces to the default value, ! 48000
host1(config)#no ip dhcp-local limit pos
```

To display the maximum number of IP address leases available for POS access interfaces, use the existing **show ip dhcp-local limits** command. For example:

```
host1#show ip dhcp-local limits

****************************************
     DHCP Local Server Address Limits
ATM Limit      - 48000
VLAN Limit     - 48000
POS Limit      - 1000
Ethernet Limit - 48000
```

Unsupported Features

The JunosE Release 11.1.x documentation set describes some features that are present in the code but that have not yet been fully qualified by Juniper Networks. If you use any of these features before they have been fully qualified, it is your responsibility to ensure that the feature operates correctly in your targeted configuration.

The following features are present but unsupported in this release.

E120 Router and E320 Router

- The ES2 10G LM and ES2 10G Uplink LM do not support layer 2 statistics for VLANs.
- Subscriber Interfaces on the ES2 10G Uplink LM

  You can configure dynamic subscriber interfaces and static subscriber interfaces on the ES2 10G Uplink LM using the CLI. However, configuring subscriber interfaces on the ES2 10G Uplink LM provides no benefit because access features such as per-subscriber QoS are unavailable on the module.

Policy Management

- External Parent Groups Unsupported on ES2 10G, ES2 10G Uplink, and ES2 10G ADV LMs

  External parent groups are not supported on the ES2 10G, ES2 10G Uplink, and ES2 10G ADV LMs. If you create a policy that references an external parent group on these LMs, the system prevents you from attaching it to the LM interface and you receive an error message.

Stateful SRP Switchover (High Availability)

- Stateful SRP Switchover for Certain Applications

  The stateful SRP switchover feature has not been qualified for the following applications:

Remote Access
– DHCP proxy client
– L2TP dialout

Release Software Protocols

The following list identifies the major software protocols supported in this release. For detailed information about any protocol, see the configuration guides.

Core Routing Stack

- Internet Protocol (IP) version 4 and version 6
- Transmission Control Protocol (TCP) for IPv4
- User Datagram Protocol (UDP) for IPv4 and IPv6

Layer 2 Protocols

- Asynchronous Transfer Mode (ATM)
- Bridged Ethernet
- Bridged IP
- Cisco High-Level Data Link Control (Cisco HDLC)
- Ethernet
- Extensible Authentication Protocol (EAP)
- Frame Relay
- Layer 2 Tunneling Protocol (L2TP)
- Multilink Frame Relay (MLFR)
- Multilink Point-to-Point Protocol (MLPPP)
- Packet over SONET (POS)
- Point-to-Point Protocol (PPP)
- PPP over Ethernet (PPPoE)
- Transparent bridging

Multiprotocol Label Switching (MPLS)

- Border Gateway Protocol (BGP-4)
- Label Distribution Protocol (LDP)
- Resource ReSerVation Protocol – Traffic Engineering Extensions (RSVP-TE)

Network Management Protocols

- Simple Network Management Protocol (SNMP) versions 1, 2c, and 3

Routing Protocols

- Border Gateway Protocol (BGP-4)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Internet Group Membership Protocol (IGMP)
- Intermediate System–to–Intermediate System (IS-IS)
- Layer 2 Virtual Private Networks (L2VPNs)
- Mobile IP
- Open Shortest Path First (OSPF) version 2 and version 3
- Protocol Independent Multicast Protocol (PIM), including PIM dense mode, PIM sparse mode, PIM dense-sparse mode, and PIM source-specific multicast
- Routing Information Protocol (RIP) version 2
- Virtual Private LAN Service (VPLS)
- Virtual Router Redundancy Protocol (VRRP)

Security Protocols

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

SRC Software and SDX Software Compatibility Matrix

The SRC software offers the features of the SDX software on the C Series Controllers, a range of hardware platforms that use the Linux operating system. In contrast, the SDX software runs on Solaris workstations. The SRC software contains the features found in the associated SDX release plus additional features described in the *SRC Release Notes*.

The following table shows which versions of the SRC software and SDX software are compatible with specified versions of the JunosE Software.

SRC Software Release	SDX Software Release	Tested with JunosE Release
2.0.0	7.1.0	8.1.2, 8.2.2
2.1.0	Not applicable	9.1.0p0-1
3.0.0	Not applicable	9.0.0, 9.0.1, 9.1.1
3.1.0	Not applicable	9.2.0, 9.3.0, 10.0.0
3.2.0	Not applicable	10.1.1, 10.2.1
4.0.0	Not applicable	10.3, 11.0, 11.1

For more detailed information about SRC software and SDX software compatibility with JunosE releases, see the *SRC Release Notes*.

Known Behavior

This section briefly describes E Series router behavior and related issues. In some cases the behavior differs from non–E Series implementations; in others the behavior is included to emphasize how the router works.

AAA

- Although you can use the **max-sessions** command to configure a maximum of 32,000 outstanding authentication/authorization requests to a RADIUS server, AAA internal limits prevent the actual number of outstanding authentication/authorization requests from exceeding 9600. These internal AAA limits apply only to authentication/authorization requests and not to accounting requests.
- The JunosE Software does not support accounting for ATM 1483 subscribers. The **atm1483** keyword for the **aaa accounting default** command is present in the CLI, but it is not supported.

ATM

- You cannot configure connection admission control (CAC) on an ATM interface on which you have created a bulk-configured virtual circuit (VC) range for use by a dynamic ATM 1483 subinterface. Conversely, you cannot create a bulk-configured VC range on an ATM interface on which you have configured CAC. The router rejects these configurations, which causes them to fail.

  Configuring CAC and bulk-configured VCs on the same ATM interface was supported in previous JunosE Software releases. As a result, If you are upgrading to the current JunosE release from a lower-numbered release, configurations that use CAC and bulk configuration on the same ATM interface continue to work. However, we recommend that you disable CAC on these ATM interfaces to ensure continued compatibility with future JunosE releases.
- When you reload an ATM line module that is configured with NBMA circuits as passive OSPF interfaces and that has established OSPF adjacencies and IBGP peers (configured on Gigabit Ethernet interfaces), the transmission of OSPF hello packets might be affected until all the NBMA interfaces have initialized.

  Work-around: Either remove the passive OSPF interface statements on the NBMA interfaces, or statically configure the OSPF cost on the NBMA interfaces.
- When you configure an ATM PVC where PCR = SCR and maximum burst size is zero, the CLI returns an error indicating the burst size is invalid and it does not create the VC.

  Work-around: Configure a CBR or a UBR plus PCR to create the circuit with the same parameters, depending on the desired priority for the traffic. CBR has a high priority and UBR plus PCR has a medium priority.
- The ATM peak cell rate (PCR) does not appear in the L2TP Calling Number AVP for the first PPP session when the ATM shaping parameters were configured by RADIUS return attributes.

- When you use the **no-authenticate** keyword with the **subscriber** command to prevent subscriber authentication so that the subscriber information can be used for DHCP option 82, suboption 2, the SRP module can reset. This issue does not occur when you use the **no-authenticate** keyword with the **subscriber** command as a way to perform a RADIUS configuration.
- Unified ISSU is not supported when ILMI is configured on ATM interfaces. [Defect ID 176007/177297/177122]
- When you perform an snmpWalk on the juniAtmSubIfVccTable, a response is received for only a few of the total configured ATM subinterfaces when both of the following are true: the router has a line module that has some ATM-related configuration and the line module is in the disabled state.

BGP

- The E Series router does not include the link-local IPv6 address in the next-hop field of an MP-BGP update message carrying IPv6 routing information over IPv4 transport. This behavior is compliant with RFC 2545 but might have interoperability issues with other implementations that depend on a link-local IPv6 address in the next-hop field on a directly connected external BGP peering.

  Work-around: Enable EBGP multihop configuration on the remote (non–Juniper Networks) peer.
- The following message might be displayed under certain conditions:

  bgpConnections (default,0.0.0.0): TCP error code xx (...) occurred while accepting inbound TCP connection

  The message is generated when an unconfigured peer attempts to establish a TCP session with an E Series router 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 E Series router; 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

BGP/MPLS VPNs

- In a scaled environment, we recommend that you increase the hold timers for the following protocols to appropriate values, based on the level of complexity of the network and scaling settings, so as to enable graceful restart to be completed successfully. [Defect ID 184974]
  - BGP
  - IS-IS
  - LDP
  - OSPF
  - RSVP
- NAT does not function properly with secondary routing table lookup (fallback global) or global export mapping on the VRF.

B-RAS

- Pool groups are not supported; although the **ip local pool group** command appears in the CLI, it is not supported.
- If the router is under a heavy load, the **show profile** command might take longer than usual to execute.

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

Bridged Ethernet

- The CLI erroneously permits you to configure **bridge1483** encapsulation over AAL5MUX IP even though that configuration is not supported.

CLI

- In Interface Configuration mode for a major interface, the CLI displays options for protocols that are not supported by that interface type.
- When you issue the **reload** command on an ERX310 router, the command might display a warning message that erroneously indicates that a synchronizing operation will be performed. Any references to synchronization that appear in command output or system messages do not apply to the ERX310 router, which does not support SRP module redundancy.
- The following commands have been deprecated in the JunosE Software and might be removed completely in a future release. If a command has been deprecated for only a particular command mode, the table specifies any modes for which it is still available.

Deprecated Command	Command Mode	Preferred Command
aaa accounting interval	Global Configuration	aaa service accounting interval and aaa user accounting interval
cablelength short	Controller Configuration	
clock rate	Interface Configuration	
channel-group description	Controller Configuration	
channel-group shutdown	Controller Configuration	
channel-group snmp trap link-status	Controller Configuration	
channel-group timeslots	Controller Configuration	
classifier-list	Global Configuration	ip classifier-list
color	Policy List Configuration	color in Classifier Group Configuration mode
controller e1	Global Configuration	
controller t1	Global Configuration	
description	Interface Configuration Still available in Controller Configuration and VRF Configuration modes	ip description
fdl	Controller Configuration	
fdl carrier	Controller Configuration	
fdl string	Controller Configuration	
fdl transmit	Controller Configuration	

Deprecated Command	Command Mode	Preferred Command
filter	Policy List Configuration	filter in Classifier Group Configuration mode
forward next-hop	Policy List Configuration	forward next-hop in Classifier Group Configuration mode
forward next-interface	Policy List Configuration	forward interface in Classifier Group Configuration mode
hostname	Domain Map Tunnel Configuration Still available in Global Configuration mode	client-name
hssi description	Interface Configuration	
hssi force dte acknowledge	Interface Configuration	
hssi internal-clock	Interface Configuration	
ignore dcd	Interface Configuration	
ignore link-state-signals	Interface Configuration	
[no] ike crl	Global Configuration	[no] ipsec crl
interface hssi	Global Configuration	
invert tx clock	Global Configuration	
ip dhcp-local cable-modem	Global Configuration	set dhcp-relay with the strings docsis and pktc in the server-string mapping specification
ip mirror	Global Configuration	ip policy secure-input and ip policy secure-output; for E120 and E320 routers, you must use these commands because the ip mirror command has been removed from the CLI for those routers.
ip policy local-input	Interface Configuration, Profile Configuration	None
[no] ipsec isakmp-policy rule	Global Configuration	[no] ipsec ike-policy-rule
ipv6 policy local-input	Interface Configuration, Profile Configuration	None
j1	Controller Configuration	
license l2tp-session	Global Configuration	None
lineCoding	Controller Configuration	
log	Policy List Configuration	log in Classifier Group Configuration mode
log severity debug dhcpLocalProtocolDecode	Global Configuration	log severity debug dhcpCapture
loopback	Domain Map Configuration Still available in Controller Configuration and Interface Configuration modes	local-interface
loopback remote { remote line fdl ansi | remote line fdl bellcore | remote line inband remote payload [fdl] [ansi] }	Controller Configuration	

Deprecated Command	Command Mode	Preferred Command
filter	Policy List Configuration	filter in Classifier Group Configuration mode
forward next-hop	Policy List Configuration	forward next-hop in Classifier Group Configuration mode
forward next-interface	Policy List Configuration	forward interface in Classifier Group Configuration mode
hostname	Domain Map Tunnel Configuration Still available in Global Configuration mode	client-name
hssi description	Interface Configuration	
hssi force dte acknowledge	Interface Configuration	
hssi internal-clock	Interface Configuration	
ignore dcd	Interface Configuration	
ignore link-state-signals	Interface Configuration	
[no] ike crl	Global Configuration	[no] ipsec crl
interface hssi	Global Configuration	
invert tx clock	Global Configuration	
ip dhcp-local cable-modem	Global Configuration	set dhcp-relay with the strings docsis and pktc in the server-string mapping specification
ip mirror	Global Configuration	ip policy secure-input and ip policy secure-output; for E120 and E320 routers, you must use these commands because the ip mirror command has been removed from the CLI for those routers.
ip policy local-input	Interface Configuration, Profile Configuration	None
[no] ipsec isakmp-policy rule	Global Configuration	[no] ipsec ike-policy-rule
ipv6 policy local-input	Interface Configuration, Profile Configuration	None
j1	Controller Configuration	
license l2tp-session	Global Configuration	None
lineCoding	Controller Configuration	
log	Policy List Configuration	log in Classifier Group Configuration mode
log severity debug dhcpLocalProtocolDecode	Global Configuration	log severity debug dhcpCapture
loopback	Domain Map Configuration Still available in Controller Configuration and Interface Configuration modes	local-interface
loopback remote { remote line fdl ansi | remote line fdl bellcore | remote line inband remote payload [fdl] [ansi] }	Controller Configuration	

Deprecated Command	Command Mode	Preferred Command
mark	Policy List Configuration	mark in Classifier Group Configuration mode
mark-de	Policy List Configuration	mark-de in Classifier Group Configuration mode
mark-exp	Policy List Configuration	mark-exp in Classifier Group Configuration mode
mark-user-priority	Policy List Configuration	mark-user-priority in Classifier Group Configuration mode
mpls ldp discovery transport-address	Interface Configuration	This command has no effect in Interface Configuration mode. Now available in Global Configuration mode.
mpls topology-driven-lsp ip-interfaces	Global Configuration	ldp ip-forwarding
[no] next-hop	Policy List Configuration	forward next-hop in Classifier Group Configuration mode
[no] next-interface	Policy List Configuration	forward interface in Classifier Group Configuration mode
nrzi-encoding	Interface Configuration	
no ospf enable	Router Configuration	ospf shutdown
policy-list	Global Configuration	ip policy-list
radius disconnect client	Global Configuration The RADIUS Disconnect Configuration mode has been removed from the CLI.	subscriber disconnect
rate-limit-profile	Policy List Configuration	rate-limit-profile in Classifier Group Configuration mode
remote-loopback	Controller Configuration	
router-name	Domain Map Configuration Still available in Tunnel Group Tunnel Configuration mode	auth-router-name and ip-router-name in Domain Map Configuration mode
show controllers t1/e1	User Exec, Privileged Exec	
show controllers t1 remote	User Exec, Privileged Exec	
show ike certificates	User Exec, Privileged Exec	show ipsec certificates
show ike configuration	User Exec, Privileged Exec	show ipsec ike-configuration
show ike identity	User Exec, Privileged Exec	show ipsec identity
show ike policy-rule	User Exec, Privileged Exec	show ipsec ike-policy-rule
show ike sa	User Exec, Privileged Exec	show ipsec ike-sa
show ip dhcp-external binding	Privileged Exec	show dhcp binding
show ip dhcp-external binding-id	Privileged Exec	show dhcp binding
show ip dhcp-local binding	Privileged Exec	show dhcp binding
show ip dynamic-interface-prefix	Privileged Exec, User Exec	None
show ip mirror interface	Privileged Exec	show secure policy-list
show license l2tp-session	User Exec, Privileged Exec	None
t1 lineCoding	Controller Configuration	None. This command never had any effect.
traffic-class	Policy List Configuration	traffic-class in Classifier Group Configuration mode

Deprecated Command	Command Mode	Preferred Command
tunnel mpls label-dist	Interface Configuration, Tunnel Profile Configuration	None
tunnel mpls autoroute announce bgp	Interface Configuration, Tunnel Profile Configuration	None
unframed	Controller Configuration	
user-packet-class	Policy List Configuration	user-packet-class in Classifier Group Configuration mode
virtual-router	Domain Map Configuration Still available in Privileged Exec and Global Configuration modes	auth-router-name and ip-router-name in Domain Map Configuration mode
yellow	Controller Configuration	

The router displays a notice when you issue the command manually. If the command is in a script, the router automatically maps the deprecated command to the preferred command. If the deprecated command no longer has a function, then that command has no effect when you run a script containing the command.

- The **show configuration** command normally takes a long time to finish for extremely large configurations. If you specify a search string (with 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.

DHCP

- Configuring authentication on the DHCP local server requires that you first disable the DHCP local server for standalone mode. Doing so removes your entire DHCP local server configuration. Therefore, if you want to configure authentication, do so before you have otherwise configured the DHCP local server.
- When you upgrade from a release numbered lower than Release 7.1.0, all DHCP host routes previously stored in NVS are deleted. After the upgrade, DHCP clients must reacquire their IP addresses, which results in the new host routes being correctly stored in NVS.

DHCP External Server

- If you are using DHCP external server and a burst of client releases occurs during a unified ISSU, some of the client releases might not be processed. [Defect ID 180178]
- When the DHCP relay agent application and the DHCP external server application are configured in the same virtual router, using the **ip dhcp-external server-sync** command on an unnumbered IP interface does not function as expected. When you issue the **ip dhcp-external server-sync** command in this configuration to create subscriber state information based on lease renewals when the external DHCP server and the router are unsynchronized, the router does not forward the ACK request from the DHCP server to the client because there is no route. [Defect ID 88562]

- When a bound DHCP client on a dynamic subscriber interface extends its IP address lease by restarting the DHCP discovery process on its primary IP interface instead of by initiating the DHCP renewal process on its dynamic subscriber interface, the default behavior of the DHCP external server application to preserve the client's dynamic subscriber interface was changed in the following JunosE releases to delete and re-create the client's dynamic subscriber interface:
  - Release 7.2.4p0-4 and all higher-numbered 7.2.x releases and patch releases
  - Release 7.3.4 and all higher-numbered 7.3.x releases and patch releases
  - Release 8.0.4 and all higher-numbered 8.0.x releases and patch releases
  - Release 8.1.2 and all higher-numbered 8.1.x releases and patch releases
  - Release 8.2.3 and all 8.2.3 patch releases
  - Release 9.0.0 and all 9.0.0 patch releases
  - Release 9.0.1 and all 9.0.1 patch releases
  - Release 9.1.0 and all 9.1.0 patch releases

  If you are upgrading the JunosE Software on the router from any of these releases, you must explicitly issue the **ip dhcp-external recreate-subscriber-interface** command to configure the router to continue to delete and re-create the DHCP client's dynamic subscriber interface.

Informational Note: The DHCP external server application is unsupported in JunosE Release 8.2.1 and JunosE Release 8.2.2.

- DHCP NAK packets are sent from a different VLAN than the one on which the renew request is received on a router that is configured with dynamic VLANs, DHCP local server, and automatically created dynamic subscriber interfaces. This behavior occurs only after a link flap has taken place.
- DHCP external server may not be able to bind all DHCP clients when all of the following conditions exist:
  - DHCP external server and either DHCP relay or relay proxy are configured in separate virtual routers on an E320 router.
  - The client-facing and server-facing interfaces for DHCP external server and either DHCP relay or relay proxy are configured on the same ES2 4G LM.
  - DHCP external server is configured to create dynamic subscriber interfaces.

  When these three conditions exist simultaneously, the ES2 4G LM may not be able to successfully process all DHCP packets. Although all clients may get bounded in DHCP relay or relay proxy, some clients may not get bounded in DHCP external server. (In a production environment it is highly unlikely for conditions 1 and 2 to exist because stand-alone DHCP external server is normally configured for a DHCP relay in a different chassis.)

Work-around: You can eliminate this issue by modifying any one of these conditions. For example, this issue does not exist with any of the following configuration modifications:

- Configure DHCP external server and either DHCP relay or relay proxy in the same virtual router.
- Configure the client-facing and server-facing interfaces for DHCP external server and either DHCP relay or relay proxy on the same ES2 10G LM instead of the same ES2 4G LM.
- Configure the client-facing and server-facing interfaces for DHCP external server and either DHCP relay or relay proxy on separate ES2 4G LMs.

Dynamic Interfaces

- Dynamic IPv6 interfaces over static PPP interfaces are not supported.

Ethernet

- The hashing algorithm that selects the LAG member link is associated with the IP address of the subscriber client to support QoS. Consequently, a particular flow is always hashed to the same link. When a member link is removed from a LAG bundle, traffic rate is disrupted and traffic flow is reduced. When the link goes down and then comes back up, the traffic flow is automatically redistributed.
- When counting bits per second on a Fast Ethernet or Gigabit Ethernet interface, an E Series router 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–E Series routers. This value therefore shows the total bandwidth utilization on the interface, including both data and overhead.
- To bridge unicast known-DA packets at line rate on both 2-Gbps ports of the GE-2 line module or the GE-HDE module when paired with the GE-2 SFP I/O module, the minimum packet size must be at least 144 bytes.

  When installed in the ERX1440 router, the GE-2 module delivers full bandwidth of 4 GB per line module (2 GB at the ingress and 2 GB at the egress) only when installed in slot 2 or slot 4, and when the SRP-40G+ module is used in the router. When installed in any other ERX1440 slot, the GE-2 module delivers a maximum bandwidth of 2 GB per line module (1 GB maximum at the ingress and 1 GB maximum at the egress). Therefore, of the maximum 24 possible ports for the module in an ERX1440 chassis (that is, two ports in each of 12 slots), full bandwidth is delivered only on a maximum of four ports (those in slots 2 and 4).

  When installed in the ERX1440 router, the GE-HDE line module delivers full bandwidth of 4 GB per line module (2 GB at the ingress and 2 GB at the egress) only when installed in slot 2 or slot 4, and when the SRP-40G+ module is used in the router. When installed in any other ERX1440 slot, the GE-HDE module delivers a maximum bandwidth of 2 GB per line module (1 GB maximum at the ingress and 1 GB maximum at the egress). Therefore, of the maximum 96 possible ports for the module in an ERX1440 chassis (that is, 8 ports in each of 12 slots), full bandwidth is delivered only on a maximum of 16 ports (those in slots 2 and 4).

  When the GE-2 line module or the GE-HDE line module is installed in either the ERX1440 router or the ERX310 router and both ports are active, line rate performance is achieved only with packets that are 174 bytes or larger. The line module might not achieve line rate with packets that are smaller than 174 bytes.

- Support for the 0x9200 S-VLAN Ethertype has been removed. You can no longer specify the **9200** option with the **svlan ethertype** command.

  When you upgrade to Release 7.1.0 or higher-numbered release, the software automatically transfers existing configurations that use the 0x9200 Ethertype to the 0x88a8 Ethertype.

- The **show interface gigabitEthernet** command output does not display the following line of output for Gigabit Ethernet modules that do not support SFPs, such as the GE Single Mode I/O module and GE I/O Multi Mode I/O modules:

```
Primary/Secondary link signal detected
Primary/Secondary link signal not detected
```

Flash

- Flash cards manufactured by Wintec are present on some currently deployed routers. When you upgrade the JunosE Software on such routers, the firmware on the flash card controller is automatically updated during diagnostics. During this reboot, the software runs an integrity check on the file system to verify that the firmware update did not corrupt the contents of the flash card. This integrity check is an expected side effect of the enhanced firmware available in this release. The integrity check does not indicate a problem with the flash card or its contents.

Forwarding

- VPLS forwarding does not function properly when any of the following conditions occur:
  - MLPPP interfaces are used
  - L2TP is used with sequence numbers enabled
  - GRE is used with sequence numbers enabled
- A memory leak of about two percent can occur on the ES2 10G LM and result in a module reset when a large number of successive SRP switchovers take place with active DHCP clients.

GRE

- When you shut down the only outgoing IP interface to the IP destinations of GRE/IP tunnels, the tunnels remain in the up state rather than transitioning to down. As a consequence, all IP routes that use these tunnels as next hops also remain in the routing table.

Hardware

- SRP modules with only 1 GB of memory do not work reliably in ERX7xx and ERX14xx routers running JunosE Release 8.1.0 or higher, and may experience system resets due to an out of memory condition. However, the ERX310 router still supports 1 GB of memory in the SRP-SE10 module.

  Work-around: Upgrade your SRP module memory to 2 GB for all ERX7xx and ERX14xx routers running JunosE Release 8.1.0 or higher.

- Do not include a **not** *protocol* clause in any classifier control list for policies attached to an interface on an ES2 10G Uplink LM. The **not** *protocol* functionality is not available for this module.
- The ES2 10G LM and the ES2 10G Uplink LM do not support VLAN statistics in the current release.

- PCMCIA NVS Card Caution

Caution: Before you insert or remove PCMCIA NVS (flash) cards from a running router, we strongly recommend that you halt the SRP module or shut down the router. Failure to do this can result in file corruption in one or both cards.

- The 4XOC3 APS MULTIMODE and 4XOC3 APS SINGLE MODE I/O modules are incompatible with the following versions of the OCx/STMx ATM and OCx/STMx POS line modules:
  - OCx/STMx ATM line modules with assembly numbers 350-00039-xx, 350-80039-xx, and 350-90039-xx
  - OCx/STMx POS line modules with assembly number 350-10039-xx
- When you configure 1:5 line module redundancy by using either the 4XOC3 APS MULTIMODE or 4XOC3 APS SINGLE MODE I/O module, the spare R-Mid OCX I/O module you install must have assembly number 350-00094-01 Rev. A01 or later. Spare R-Mid OCX I/O modules with an earlier assembly number are not supported for 1:5 redundancy configurations that use either the 4XOC3 APS MULTIMODE or 4XOC3 APS SINGLE MODE I/O module.
- There is a very small chance that some line modules can have an improperly modified keying block that prevents the module from proper seating in the top slot of an older ERX7xx chassis or a preproduction ERX310 chassis. For example, this problem has been observed for an OCx/STMx module in slot 2 of an early-test ERX310 chassis.

  Work-around: Remove the keying block to insert the module into the top slot, or insert the module into a different slot.

HDLC

- By design, on the cOC12/STM4 module you cannot delete a serial interface while data for the interface is still enqueued. The enqueued data can drain only when the interface is operationally up. Therefore you must ensure that the interface is operationally up before you delete it. For example, if you have issued the **shutdown** command for the interface before you try to delete the interface, issue the **no shutdown** command, then delete the interface.

IP

- The **ip tcp adjust-mss** command, which modifies the maximum segment size for TCP SYN packets traveling through the interface, is not supported on the ES2 10G LM or ES2 10G Uplink LM.
- If you have enabled ipInterface logging at a priority of debug, the acknowledgment 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. This behavior has no functional effect other than that the log is misleading. However, previous log events indicate that the interface deletion was beginning.

- When you upgrade from certain releases to JunosE Release 9.2.0p1-0 or higher-numbered releases, descriptions configured for IP interfaces or IP subinterfaces are not retained across the upgrade when the descriptions are shorter than 9 characters in length. Additionally, VRF descriptions are not retained across the upgrade when the combined length of the VRF description and the VRF name is shorter than 9 characters. This behavior is seen during upgrades using a reload, stateful SRP switchover, or unified ISSU. Upgrades from the following releases are affected by this behavior:
  - 7.x.x
  - 8.0.x
  - 8.1.x, 8.2.x, and 9.x.x builds created before July 23, 2008

  Examples of descriptions that are not retained across the upgrade:

  host1(config-if)#**ip description 12345678**

  host1(config)#**ip vrf 123**
  host1(config-vrf)#**description 45678**

  Examples of descriptions that are retained across the upgrade:

  host1(config-if)#**ip description longdescription**

  host1(config)#**ip vrf longername**
  host1(config-vrf)#**description 45678**

  host1(config)#**ip vrf 123**
  host1(config-vrf)#**description longdescription**

  Work-around: Before you upgrade from an affected release to JunosE Release 9.2.0p1-0 or higher-numbered releases, ensure that you do the following:
  - Change IP interface and subinterface descriptions to 9 or more characters.
  - Change VRF descriptions, VRF names, or both so that the combination of associated VRF names and descriptions consists of 9 or more characters.
- When you want to use a configuration script to configure IP shared interfaces that reference a physical interface, you must issue the **service show configuration format 2** command before you generate the script. If the default **show configuration** format (format 1) is enabled instead, the generated script cannot properly configure the IP shared interfaces because they are created before the physical interfaces. To properly configure the shared interfaces in this event, run the generated format 1 script twice.
- When you issue the **show ip forwarding-table** command for a particular slot, it is normal and appropriate behavior when 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 indicates Invalid, then the routing table distribution has failed constantly for that VR and a real problem exists. You might occasionally see a status of Updating. However, if the Status field always indicates Updating, then again the routing table distribution has failed constantly for that VR, and a real problem exists.

- 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.
- The enhancement to the CLI to support unnumbered reference to any kind of interface rather than just loopback interfaces has consequences such as the following: [Defect ID 47743]
  - If the references to shared interfaces appear in the **show configuration** output before the configuration for the interfaces they refer to, trying to restore such a configuration with a script generated from **show configuration** generates errors like the following:

    ```
    % Error, line 3929:
    host1(config-if)#ip share-interface FastEthernet 3/0.2
    % No such interface
    ```

  - Unnumbered interfaces that refer to nonloopback interfaces (for example, **ip unnumbered fastEthernet 3/0.2**) and that appear in the **show configuration** output before the interface referred to might generate similar no such interface errors.

  Work-around: Run the script twice.
- If you enable detection of duplicate IPv6 prefixes using the **aaa duplicate-prefix-check** command, and bring up a subscriber in a dual-stack network (in which both IPv4 and IPv6 subscribers are present) over a static PPP interface for which IPv6 prefix is configured for IPv6 Neighbor Discovery router advertisements (using the **ipv6 nd prefix-advertisement ipv6Prefix** command), the subscriber session is successfully brought up. When you attempt to bring up another subscriber over a different interface on the same virtual router as the one used for the first subscriber, and for which the Ipv6-NdRa-Prefix (VSA 26-129) returned from the RADIUS server in the Access-Accept message is the same IPv6 prefix as the statically configured value for the first subscriber, the second subscriber session is also brought up and not disconnected as expected.

  In such a scenario, the duplicate IPv6 prefix detection functionality does not cause the second subscriber session, which uses the same IPv6 prefix as the first subscriber session, to be rejected. Also, a new IPv6 route is installed for the second subscriber as a duplicate access-internal route. [Defect ID 187264]

IPSec

- When you shut down the only outgoing IP interface to the IP destinations of IPSec tunnels, the tunnels remain in the up state rather than transitioning to down. As a consequence, all IP routes that use these tunnels as next hops also remain in the routing table. You can use dead keepalive detection (DPD) to avoid this situation. DPD must be active, which requires both IPSec tunnel endpoints to support DPD.

- During a warm restart after a system failover, the SRP module can take several minutes to resume the normal exchange of UDP/IP packets to applications. During this restart time, the E Series router does not send or receive dead peer detection (DPD) keepalives, which are used to verify connectivity between the router and its peers. The length of the restart time depends on the number of interfaces—if the restart time is too long, remote peers might determine that the connection from them to the E Series router is broken and then shut down an IPSec tunnel that has DPD enabled. In the worst case, all IPSec tunnels might be shut down. [Defect ID 65132]
- When the LAC–to–LNS data path runs over an MPLS tunnel and the MPLS tunnel originates or terminates at the LAC on an ES2 10G LM or an ES2 10G Uplink LM, the L2TP data traffic that originated or terminated at the LAC is discarded.

IS-IS

- When IS-IS is configured on a static PPP interface, the IS-IS neighbor does not come up if you remove the IP address from the interface and then add the IP address back to the interface.

  Work-around: When you remove and add back the IP address, you must also remove the IS-IS configuration from the interface and then add the configuration back to the interface by issuing the **no router isis** and **router isis** commands.
- When you run IS-IS on back-to-back virtual routers (VRs) in an IS-IS-over-bridged-Ethernet configuration and do not configure different IS-IS priority levels on each VR, a situation can occur in which both VRs elect themselves as the designated intermediate system (DIS) for the same network segment.

  This situation occurs because the router uses the same MAC address on all bridged Ethernet interfaces by default. When both VRs have the same (that is, the default) IS-IS priority level, the router must use the MAC address assigned to each interface to determine which router becomes the DIS. Because each interface in an IS-IS-over-bridged-Ethernet configuration uses the same MAC address, however, the router cannot properly designate the DIS for the network segment. As a result, both VRs elect themselves as the DIS for the same network segment, and the configuration fails. [Defect ID 72367]

  Work-around: To ensure proper election of the DIS when you configure IS-IS over bridged Ethernet for back-to-back VRs, we recommend that you use the **isis network point-to-point** command in Interface Configuration mode to configure IS-IS to operate using point-to-point (P2P) connections on a broadcast circuit when only two routers (or, in this case, two VRs) are on the circuit. Issuing this command tears down the current existing IS-IS adjacency in that link and reestablishes a new adjacency.

- If you perform a stateful SRP switchover operation on a router with IS-IS previously configured on the device, the IS-IS application takes longer than the normal duration (approximately 40 seconds) to restart after the switchover is completed. The time that it takes for IS-IS to restart after a stateful switchover causes a large delay in the transmission of hello packets with restart TLV (type 211) from the restarting router to neighboring routers. Because of the delay in transmission of hello packets to neighboring routers, active adjacencies are not maintained between the restating router and other routers in the IS-IS domain. To avoid adjacencies being reset, we recommend that you increase the hold timers for the IS-IS protocol to appropriate values, based on the level of complexity of the network and configuration settings, so as to enable IS-IS graceful restart to be completed successfully.

  The long duration for restart of a previously running application on the router also occurs if you configured OSPF on the router and perform a stateful SRP switchover process. This condition can occur even in environments that are not scaled to the maximum limits and contain minimal subscriber connections or attribute definitions.

  Because the IP application takes about 30-35 seconds to reinitialize and process control packets after a stateful SRP switchover, and the continual increase in the time needed for completion of IP reinitialization in JunosE releases (owing to consumption of system resources for enhanced functionalities), we recommend that you increase the hold timers for the associated protocols running on the router to necessary levels so that graceful restart can function properly. [Defect ID 90546]

L2TP

- NAT dynamic translation generation affects the LNS session creation time. When NAT dynamic translations and LNS sessions are created simultaneously, NAT dominates the CPU cycles of the tunnel-service module, resulting in a delay in the LNS session creation rate. The LNS session creation rate returns to its normal rate when NAT dynamic translations are no longer being generated. [Defect ID 53191]

  Work-around: When signaling performance must be optimal, avoid the simultaneous configuration of NAT and LNS.

- L2TP peer resynchronization enables an L2TP failed endpoint to resynchronize with its peer non-failed endpoint. The JunosE Software supports failover protocol and silent failover peer resynchronization methods. If you configure the silent failover method, you must keep the following considerations in mind:

  – PPP keepalives—To ensure resynchronization of the session database, PPP keepalives must be enabled on the L2TP data path. Without PPP keepalives, silent failover might disconnect an established session if there is no user traffic during failover recovery.

  – Asymmetric routes on different line modules—Asymmetric routes whose receive and transmit paths use I/O paths on different line modules can result in improperly handled line module control packets. If your network does include this type of asymmetric route, tunnels using these routes might fail to recover properly.

- 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 E Series router 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, use a different name for the destination profile. [Defect ID 32973]

Line Module Redundancy

- On E120 routers and E320 routers, redundant IOAs have a temperature sensor, and the **show environment all** command lists the temperature of IOAs in their associated slots.

  On ERX routers, redundant I/O modules do not have a temperature sensor. Therefore, the **show environment all** command output lists the primary I/O module temperature in the slot of the line module that is responsible for the I/O module.

- When you install an ES2-S1 Redundancy IOA with a hardware revision number of -02 or less in slot 0 or slot 11 of the E320 router or in slot 0 or the E120 router, do not install an OCx/STMx ATM IOA or an OCx/STMx POS IOA in the lower (E320) or left (E120) adapter bay of slot 1 or slot 12. When the spare line module is controlling another slot and you revert back to the primary line module, the ATM or POS IOAs can become unusable or cause the line module to reset. [Defect ID 69760]

  Work-around: This problem is not present for ES2-S1 Redundancy IOAs with a hardware revision number of -03 or higher.

MLPPP

- Do not configure both MLPPP fragmentation (with the **ppp fragmentation** command) and IP fragmentation of L2TP packets (with the **ip mtu** command) on the same interface. Instead, you must choose only one of the fragmentation configurations by setting it to the necessary value and set the other fragmentation configuration to the maximum allowable value.

MPLS

- Martini circuits configured on the ES2 10G LM act as true layer 2 tunnels, without modifying the layer 2 headers. For this reason, Martini VLAN retagging is not currently supported.
- If you are upgrading to Release 7.1.0 or a higher-numbered release from a release numbered lower than Release 7.1.0, and have inter-AS option B or C configurations, you must explicitly configure MPLS on all inter-AS links, as in the following example:

  host1#**configure terminal**
  host1(config)#**interface fastEthernet 2/0**
  host1(config-if)#**ip address ...**
  host1(config-if)#**mpls**

  If you do not explicitly configure MPLS on the links, the inter-AS feature will not work properly.

- You cannot use an underscore character (_) in an MPLS tunnel name.

- When you upgrade the router to JunosE Release 7.1.0 or a higher-numbered release from a release numbered lower than Release 7.1.0, remote ATM layer 2 over MPLS circuits (also known as MPLS shim interfaces) that use Martini encapsulation are erroneously signaled with the control word attribute setting “Control word is not preferred by default”. Because control words are required for these MPLS shim interfaces, these circuits should instead be signaled with the setting “Control word is preferred by default”.

  Work-around: To reinstate the proper setting (“Control word is preferred by default”), remove the MPLS shim interface from the ATM subinterface and then reconfigure it.

Multicast

- Do not configure a multicast group with more than 10,219 outgoing interfaces (OIFS) on the same ES2 10G LM. [Defect ID 81768]
- When you upgrade a router running a release earlier than JunosE Release 8.2.x to JunosE Release 8.2.x or higher-numbered releases, the Protocol Independent Multicast (PIM) configuration settings in VPN routing and forwarding (VRF) instances are not restored after the upgrade is completed. This problem happens only if you did not previously configure PIM on the parent virtual router (VR) for the VRF. This problem occurs with both IPv4 PIM and IPv6 PIM configurations on the router.

  After the completion of the upgrade process, if you attempt to restore the PIM configuration directly on the VRF, an error message is displayed. For example, if you try to restore the IPv4 PIM settings on the VRF using the **router pim** command, the following error message is displayed:

```
host1:vrf01(config)#router pim
% PimIp not configured on this router
```

  Work-around: To correct this problem after you upgrade a router running a release earlier than JunosE Release 8.2.x to JunosE Release 8.2.x or higher-numbered releases, you need to restore the PIM configuration on the upgraded router in two steps (first, on the parent VR, and then, on the VRF), instead of attempting to restore the PIM configuration directly on the VRF.

  To restore IPv4 PIM configuration on the VRF, perform the following steps. These steps assume that a parent VR context, named “parent”, and a VRF in the parent VR, named “vrf01”, are already configured on the router.

  1. Access the context of the parent VR, and create and enable IPv4 PIM on the parent VR.

     host1(config)#**virtual-router parent**
     host1:parent(config)#**router pim**

  2. Enter the VRF Configuration mode to restore PIM settings on the VRF in the parent VR.

     host1:parent(config)#**virtual-router parent:vrf01**

  3. Create and enable IPv4 PIM on the VRF in the parent VR.

     host1:parent:vrf01(config)#**router pim**

After the IPv4 PIM configuration is recovered on the VRF, you can remove the IPv4 PIM configuration settings on the parent VR by using the **no router pim** command, if necessary.

To restore IPv6 PIM configuration on the VRF, perform the following steps. These steps assume that a parent VR context, named "parent", and a VRF in the parent VR, named "vrf01", are already configured on the router.

1. Access the context of the parent VR, and create and enable IPv6 PIM on the parent VR.

   host1(config)#**virtual-router parent**
   host1:parent(config)#**ipv6 router pim**

2. Enter the VRF Configuration mode to restore PIM settings on the VRF in the parent VR.

   host1:parent(config)#**virtual-router parent:vrf01**

3. Create and enable IPv6 PIM on the VRF in the parent VR.

   host1:parent:vrf01(config)#**ipv6 router pim**

After the IPv6 PIM configuration is recovered on the VRF, you can remove the IPv6 PIM configuration settings on the parent VR by using the **no ipv6 router pim** command, if necessary.

- The **ip dipe sg-cache-miss** and **ipv6 dipe** commands are not intended or supported for customer use, although they are visible in the User Exec and Privileged Exec modes respectively. These commands are intended to be used in a Juniper Networks internal lab environment for testing without a traffic generator.

Packet Mirroring

- When both interface-specific mirroring and user-specific mirroring are configured on the same interface, the interface-specific secure policies take precedence. The interface-specific secure policies, which you manually attach using the CLI, override and remove any existing secure polices that were attached by a trigger action. If the interface-specific secure polices are subsequently deleted, the original trigger-based secure policies are not restored.
- Typically, when configuring packet mirroring, you configure a static route to reach the analyzer device through the analyzer port. If the analyzer port is an IP-over-Ethernet interface, you must also configure a static Address Resolution Protocol (ARP) entry to reach the analyzer device. However, because only a single static ARP entry can be installed for a given address at any given time, when you are using equal-cost multipath (ECMP) links to connect to the analyzer device, the static ARP configuration does not provide failover if the link being selected fails or is disconnected. Therefore, to provide continued connectivity if the link fails when using ECMP, enable the **ip proxy-arp unrestricted** command on the next-hop router for each ECMP interface. As a result, when the link fails, the router sends an ARP request to identify the MAC address of the analyzer device and gets a response over the new link.
- The ES2 10G LM supports the packet mirroring feature when the module is paired with the ES2-S2 10GE PR IOA, the ES2-S1 GE-8 IOA, or the ES2-S3 GE-20 IOA. When you use the ES2 10G LM with these IOAs, CLI-based interface-specific mirroring is not supported.

Policy Management

- Multiple Forwarding Solution Rules for a Single Classifier List in a Policy

  Before Release 5.2.0, it was possible to configure a policy with multiple rules that specified forwarding solutions where all of these rules were associated with a single classifier list. This typically was a configuration error, but the CLI accepted it. Beginning with Release 5.2.0, the CLI no longer accepts this configuration.

  – Multiple forwarding rules behavior for releases numbered lower than Release 5.2.0:

    > If multiple forward or filter rules were configured to reference the same classifier list in a single policy, then all rules except the first rule configured were marked as eclipsed in the **show policy** command display. Next-interface and next-hop rules were treated in the same manner. The eclipsed rules were not applied.

    > If a policy were configured with one rule from the [forward, filter] pair and one rule from the [next-hop, next-interface] pair, and if both rules referenced the same classifier list, then no visible eclipsed marking occurred. However, these two rules were mutually exclusive, and only one of them defined the forwarding behavior. The rule action that was applied was in the order (from highest to lowest preference): next interface, filter, next hop, forward. The applied rule was the rule whose behavior was seen by forwarded packets.

      For example, if a policy had both a next-interface and a filter rule, then the next interface was applied. If a policy had a next-hop and a filter rule, then the filter rule was applied.

  – Multiple forwarding rules behavior for Release 5.2.0 and higher-numbered releases:

    Beginning with Release 5.2.0, the multiple rules behavior is designed so that when a forwarding solution conflict occurs within a policy, such as those described earlier, the second forwarding solution overwrites the preceding solution. That is, the last forwarding rule configured for the given classifier list within a policy is the forwarding behavior that is used. Also, a warning message is now displayed when this type of conflict occurs.

    Example 1—In this example, the filter rule action overwrites the forward rule, and is therefore applied.

```
host1(config)#policy-list wstPolicyList
host1(config-policy-list)#forward classifier-group svaleClacl1
host1(config-policy-list)#filter classifier-group svaleClacl1
WARNING: This rule has replaced a previously configured rule.
host1(config-policy-list)#exit
host1(config)#
```

Example 2—In this example, three forwarding solution conflicts result in rules being overwritten. The filter rule is the last rule configured, and is therefore applied.

```
host1(config)#policy-list bostTwo
host1(config-policy-list)#forward classifier-group clacl5
host1(config-policy-list)#next-hop 1.1.1.1 classifier-group clacl5
WARNING: This rule has replaced a previously configured rule.
host1(config-policy-list)#next-interface atm 1/0.0 classifier-group clacl5
WARNING: This rule has replaced a previously configured rule.
host1(config-policy-list)#filter classifier-group clacl5
WARNING: This rule has replaced a previously configured rule.
host1(config-policy-list)#exit
host1(config)#
```

Informational Note: When you upgrade the nonvolatile memory to Release 5.2.0 or later, the upgrade removes eclipsed rules and rules whose behavior was not applied in the previous release. This removal ensures that the postupgrade forwarding behavior is the same as the preupgrade behavior.

Informational Note: If you upgrade to Release 5.2.0 or later and then configure your router using a script generated before Release 5.2.0, the postupgrade and preupgrade forwarding behaviors might not be the same. The new Release 5.2.0 configuration behavior is applied—the last policy rule configured for a given classifier list that specifies a forwarding behavior is the only rule remaining.

- You cannot configure classifier lists that reference multiple fields for a VLAN policy list on the ES2 10G Uplink LM or the ES2 10G LM, with the exception of traffic-class and color. The system incorrectly classifies VLAN policies that classify using multiple fields. For example, an invalid policy list that references multiple fields uses both color and user-packet-class, or one classifier list using color and another using user-packet-class.
- In JunosE Release 11.0.0 and higher-numbered releases, you must specify at least one option by which the router defines a packet flow in order to configure classifier control lists (CLACLs) for policy lists to be attached to VLAN interfaces. Although a carriage return, **<cr>**, is displayed when you type a question mark (?) after entering the **vlan classifier list** *classifierName* command without defining any other keyword or CLACL option, an error message is displayed when you press **Enter** to configure the VLAN CLACL with only the name. The error message states that a VLAN classifier list cannot be configured without any classification criteria, such as color, traffic class, user packet class, or user priority. You must specify at least one keyword or option to configure VLAN CLACL successfully. [Defect ID 184139].

  In JunosE releases earlier than Release 11.0.0, you could configure all CLACLs (except those CLACLs that were attached to IP interfaces) without specifying an option or a keyword. Because the policy management application treats only one default classifier group (configured with an * in the policy list) as a valid setting, this functionality change ensures that only one classifier that matches all packets can be present in a VLAN policy list definition.

- When you attach a policy to an interface and the policy contains a classifier rule that is unsupported for that interface, the CLI generates a message and the policy is applied. However, if an existing policy is already attached to that interface, then support for the new policy is not checked and the invalid policy is applied to the interface without warning. The results of this attachment are not predictable.
- The ES2 10G LM does not support the **deprecated next-hop** command.
- In rare cases, some policy configurations that use CAM hardware classifiers from releases earlier than Release 7.1.0 can fail because they exceed the total hardware classifier entry size of 128 bits that was introduced in Release 7.1.0. For more information and examples of previous configurations, see *JunosE Policy Management Configuration Guide, Chapter 8, Policy Resources*.
- No logs are created if you use the **policy-list** option with the **log severity** *severityValue* **policyMgrPacketLog policy-list** *policyListName* command when logging policyMgrPacketLog events.

PPP

- The GE-2 line module does not support dynamic IP interfaces over static PPP interfaces when the PPPoE subinterface is also static. The OC3/STM1 GE/FE line module does not support dynamic IP interfaces over static PPP interfaces when the ATM interface column is also static.

PPPoE

- On the ES2 4G LM, ES2 10G LM, and ES2 10G Uplink LM, data packets for PPPoE are not counted at the PPPoE interface. Instead, PPPoE data packets are counted at the PPP interface that sits on the PPPoE interface. Use the **show ppp interface** command to display the data packets. Control packets for PPPoE are counted at the PPPoE interface; use the **show pppoe interface** command to display the control packets.

QoS

- In JunosE Releases 7.1.x, 7.2.x, and 7.3.x, you can attach a QoS profile to Ethernet interfaces that are configured in a link aggregation group (LAG) interface. However, beginning with JunosE Release 8.0.1, you can attach a QoS profile directly to the LAG interface. As of JunosE Release 8.0.1, the software restricts you from attaching a QoS profile to any Ethernet interfaces that are members of a LAG. [Defect ID 84632]

  Work-around: Prior to upgrading from JunosE Releases 7.1.x, 7.2.x, or 7.3.x to JunosE Release 8.0.x or higher-numbered releases, remove the QoS profile from the Ethernet interface. When you have successfully upgraded to JunosE Release 8.0.x or higher-numbered releases, reattach the QoS profile to the LAG interface.
- The router cannot resolve inconsistent requests caused by two QoS profiles that modify the same scheduler property inconsistently.

  Work-around: Avoid using two QoS profiles that modify the same scheduler property inconsistently, such as setting different values for the shaping rate for the same S-VLAN node.
- When you perform an SNMP walk of the juniQosQueueStatistics MIB, a timeout of up to 5 minutes ensues, during which the SRP module CPU utilization goes to 100 percent.

By default, simple implicit shared shapers activate scheduler nodes in named traffic-class groups. The implicit constituent selection process is now the same for both simple and compound shared shapers.

This is a change in default behavior. For releases before Release 7.2.0, you could not configure scheduler nodes as active constituents of the simple shared shaper, except for the best-effort node.

To recover the default behavior available before Release 7.2.0, or to select active constituents that are different, use simple explicit shared shapers to select best-effort nodes only.

- When you are configuring compound shared shaping using explicit constituents and you explicitly specify both a scheduler node and a queue stacked above the node as constituents of the shared shaper, the system selects the scheduler node (but not the queue) as the constituent.
- Egress strict-priority packets may experience high latency on OC3/STM1 ATM interfaces associated with the LM if you have shaped the port rate to more than 148.5 Mbps.

  Work-around: To ensure low strict-priority latency, shape the port rate to no more than 148.5 Mbps.
- Egress traffic may be dropped on OC12/STM4 ATM interfaces if you have shaped the port rate to more than 542 Mbps.

  Work-around: Do not exceed a shaped port rate of 542 Mbps.
- Incorrect output is sent to the CLI the first time you enter Global Configuration mode or issue the **show subscribers** command after viewing the VLAN subinterface over which a subscriber is connected.
- An error message regarding the qos-parameter instance QosParameterDefinition is erroneously generated on an ERX1440 router when it is configured for L2C and QoS RAM and receives TLV 144 (DSL Type). The parameter instantiation actually functions properly.
- On the E120 and E320 routers, you cannot attach QoS profiles to L2TP tunnels by means of the CLI because the CLI does not pass the router ID to QoS.

RADIUS

- JunosE Software provides extended commands for configuring the formats of the RADIUS NAS-Port attribute (attribute 5) and the RADIUS Calling-Station-ID attribute (attribute 31) when the physical port value is greater than 7.

  When the physical port value is greater than 7:

  – An incorrectly configured NAS-Port attribute format results if you use either the **radius nas-port-format 0ssssppp** or **radius nas-port-format ssss0ppp** command.

  – An incorrectly configured Calling-Station-ID attribute results if you use either the **radius calling-station-format fixed-format** command or the **radius calling-station-format fixed-format-adapter-embedded** command.

Work-around: Use the following commands on routers that have line modules with more than 7 physical ports:

- To configure the NAS-Port attribute format, use the **radius nas-port-format extended [atm | ethernet]** command.
- To configure the Calling-Station-ID attribute format, use the **radius calling-station-format fixed-format-adapter-new-field** command.

SNMP

- SNMP MIBs

  Information about all the SNMP MIBs (both standard and proprietary) that the router supports in this release is available in the MIB directory in the SW_Image_CD-2 folder of the JunosE Software image bundle, which you downloaded from the Juniper Networks website, that contains the release file for E120 and E320 routers.

- Some Juniper Networks SNMPv1-formatted traps contain an incorrect object identifier (OID) in the SNMPv1-Trap-PDU enterprise field. An SNMPv2 trap is typically identified by an OID that ends in the form ...x.y.z.0.n. This OID appears, in full, as the value of the snmpTrapOID.0 object in the varbind list of an SNMPv2-formatted trap. In the corresponding SNMPv1-formatted trap, this OID is broken down into subcomponents that fill the SNMPv1-Trap-PDU enterprise field (...x.y.z) and specific trap number field (n); the zero is unused.

  The SNMPv1-formatted versions of the following Juniper Networks traps incorrectly contain ...x.y.z.0 in the SNMPv1-Trap-PDU enterprise field. That is, a zero is mistakenly appended to the correct enterprise OID value.

Trap Name	Expected Enterprise OID	Enterprise OID Sent by SNMP Agent
junidApsEventSwitchover	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
junidApsEventModeMismatch	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
junidApsEventChannelMismatch	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
junidApsEventPSBF	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
junidApsEventFEPLF	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
juniAddressPoolHighAddrUtil	.1.3.6.1.4.1.4874.2.2.21.3	.1.3.6.1.4.1.4874.2.2.21.3.0
juniAddressPoolAbatedAddrUtil	.1.3.6.1.4.1.4874.2.2.21.3	.1.3.6.1.4.1.4874.2.2.21.3.0
juniAddressPoolNoAddresses	.1.3.6.1.4.1.4874.2.2.21.3	.1.3.6.1.4.1.4874.2.2.21.3.0
juniDhcpLocalServerPoolHighAddrUtil	.1.3.6.1.4.1.4874.2.2.22.3	.1.3.6.1.4.1.4874.2.2.22.3.0
juniDhcpLocalServerPoolAbatedAddrUtil	.1.3.6.1.4.1.4874.2.2.22.3	.1.3.6.1.4.1.4874.2.2.22.3.0
juniDhcpLocalServerPoolNoAddresses	.1.3.6.1.4.1.4874.2.2.22.3	.1.3.6.1.4.1.4874.2.2.22.3.0
pimNeighborLoss	.1.3.6.1.3.61.1	.1.3.6.1.3.61.1.0

Work-around: Use the OIDs that the SNMP agent sends.

- When you configure the router with an address pool that has two IP address ranges, only the range that you configured first is available via the MIB.

SRC Software and SDX Software

- The SRC client does not prevent you from changing the name of the router while the client is connected to the SAE, resulting in SAE issues such as lost IP addresses and stale users.

  Work-around: To change the router name while the SRC client is connected to the SAE, shut down the SRC client, change the name, then re-enable the SRC client.

SSH

- If the SRP module restarts when SSH is configured in a VR other than default, SSH can sometimes become disabled. This happens if SSH attempts to bind with a VR before the VR comes back up after the restart. In this event, a warning message is generated to alert you to the fact that SSH is disabled in that VR. You must manually re-enable SSH either by accessing the console VTY or creating a Telnet session to the router.

Stateful SRP Switchover (High Availability)

- Additional processing is required to maintain and mirror the necessary state information that enables subscriber sessions to stay up across an SRP failover. As a result, the performance of other control plane functions is reduced. Specifically, call setup rates are lower than in previous releases.

Informational Note: Rapid call setup rates are most important following an outage that causes all subscribers to drop, because many of the dropped subscribers will immediately attempt to reconnect. This type of outage occurs far less frequently with stateful SRP switchover.

  We have ongoing development activities to characterize and improve call setup rates in future releases.

- Stateful SRP switchover remains inactive for 20 minutes after an initial cold-start or cold-restart of the router. This delay enables the system to reach a stable configuration before starting stateful SRP switchover.

  If you want to override the 20-minute timer, turn high availability off by using the **mode file-system-synchronization** command, and then on again by using the **mode high-availability** command.

- When IP tunnels are configured on a router enabled for stateful SRP switchover, and the Service Module (SM) carrying these tunnels is reloaded, stateful SRP switchover transitions to the pending state. Stateful SRP switchover remains in the pending state for 10 minutes following the successful reloading of the SM. This amount of time allows for IP tunnel relocation and for the tunnels to become operational again on the SM. If an SRP switchover occurs while in the pending state, the router performs a cold restart.

  Work-around: None.

- After a stateful SRP switchover, each layer of the interface columns must reconstruct its interfaces from the mirrored information. While the interfaces are being reconstructed the SRP module cannot send or receive frames, including the protocol frames that signal graceful restart behavior with OSPF and IS-IS peers. If the configured hold time is too short, peers might mistakenly declare the adjacency down during the time in which the graceful restart is taking place. [Defect ID 65132]

Work-around: Increase the hold time to provide sufficient time for interface synchronization before the peers declare the adjacency down.

- For OSPF, use the **ip ospf dead-interval** command to set the hold time. We recommend that you use Bidirectional Forwarding Detection (BFD) with a longer OSPF dead interval to achieve fast failure detection.
- For IS-IS, use the **isis hello-interval** and **isis hello-multiplier** commands to set the hold time.

We recommend the following hold times for each protocol, based on the number of interfaces.

Interface Count	Recommended Hold Time for OSPF	Recommended Hold Time for IS-IS
16000 or less	80 seconds	50 seconds
16001 to 32000	87 seconds	55 seconds
32001 to 48000	90 seconds	70 seconds

- When you issue **show** commands as soon as the CLI is available after a stateful SRP switchover, the commands can hang until the warm restart is completed.

Subscriber Interfaces

- MAC address validation is not supported on either of the following:
  - Packet-triggered subscriber interfaces that are created dynamically
  - Packet-triggered subscriber interfaces that are managed on the primary IP interface

  A packet-triggered subscriber interface is created when the router receives a packet with an IP source address that does not match any entries in the demultiplexer table. When the router detects an unmatched packet, it generates a trigger event that determines whether to create a dynamic subscriber interface or configure an existing interface. To configure packet detection on the router, use the **ip auto-detect ip-subscriber** command.

System

- ERX routers display different behavior from E120 routers and E320 routers when reporting modules as inactive.

  ERX routers report a module as inactive when either:
  - The I/O module is not present
  - The primary line module is fully booted and ready to resume operation. In this case, the standby is currently providing services.

  E120 routers and E320 routers report a module as inactive when either:
  - The primary line module has no IOAs.
  - The primary line module has IOAs, but they have failed diagnostics.
  - The standby line module has taken over for the primary line module, and has control of the IOAs.

Because E120 and E320 routers can accommodate up to two IOAs per slot, at least one IOA must be online. If the second IOA fails, the line module is still online, but does not use both IOAs. You can ensure that every module is up and active in the system and not in a failed state by issuing the **show version all** command.

- 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 an ERX1410 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 E Series router verifies whether it has available space equal to at least twice the size of the .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 .rel or .cnf files.
- 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 re-create it after deletion. [Defect ID 34842/45063]

  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.

System Logging

- 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 because SNMP polling on line modules (correctly) detects no fabric:

  NOTICE 09/01/2002 18:47:52 CEST controlNetwork (slot 11): Control Bus Master slave error 0x5 while accessing slot
- The **show configuration category management syslog virtual-router default** command incorrectly displays logs for multiple syslog destinations when you add a log to only one syslog destination. The **show log configuration** command shows the correct configuration.

Tunneling

- When you configure the GE-2 line module, the GE-HDE line module, or the ES2-S1 GE-4 IOA to operate as a shared tunnel-server module, the available bandwidth for tunnel services is limited to 0.5 Gbps per module.
- In releases numbered lower than Release 7.3.0, a dynamic tunnel-server port was located on port 8 of the GE-HDE line module and GE-8 I/O module.

  In Release 7.3.0 and higher-numbered releases, the dynamic tunnel-server port is located on port 9. When you upgrade to Release 7.3.0, any existing tunnel-server port configurations move from port 8 to port 9.

Known Problems and Limitations

This section identifies the known problems and limitations in this release. For more information about known problems that were discovered at customer sites, you can log in to the JunosE Knowledge Base at https://www2.juniper.net/kb/, enter the defect ID number in the Search by Keyword field, and click Search.

ANCP

- On an E320 router that has established 3000 ANCP adjacencies with a client and traffic is initiated, the following behavior occurs sporadically: All existing Telnet sessions are disconnected and no new Telnet sessions can be established for several minutes. [Defect ID 83872]

ATM

- The line module resets when you issue the **show nbma arp** command after you have configured NBMA interfaces on an ATM line module. [Defect ID 88491]
- When 16,000 PPPoA interfaces are configured on an OCx/STMx ATM line module paired with an OC3-4 I/O module in an ERX14xx model, ERX7xx model, or ERX310 router, Ping traffic passing through the line module on the restarting router experiences an outage of 103 seconds, which is beyond the maximum limit, after a unified ISSU from JunosE Release 9.2.0p1-0 to 9.3.0b0-12. This outage does not occur when the same configuration is applied on a Gigabit Ethernet interface. [Defect ID 179794]
- When a mirror rule that triggers on username is employed for packet mirroring of dynamic IP subscribers over ATM, removal of the rule does not disable packet mirroring. [Defect ID 175356]

  Work-around: Use a mirror rule that triggers on account session ID rather than on username.
- When you issue the **no atm atm1483 auto-configure upperInterfaceType lockout-time** command in Profile Configuration mode, the lockout time range does not revert to the default values. [Defect ID 66544]
- When one or more ATM1483 attributes appears in a profile, the **show configuration include-defaults** command fails to display the default values for all possible ATM1483 attributes. [Defect ID 67157]
- The output of the **show atm arp** command displays only 4096 entries when the line module is configured with more than 4096 NBMA ARP entries. [Defect ID 68849]
- The **baseline interface atm** command fails for a VCD assigned by the router to F4 OAM circuits. [Defect ID 174482]
- ATM line modules reset after unified ISSU completes at the LAC when an MLPPP bundle with three links are tunneled to the LNS. [Defect ID 178821]
- For PPPoE, the AAL5 inPacket Discards counter might increment erroneously during call setup when a packet is passed directly to PPPoE for negotiation rather than being discarded. [Defect ID 51757]

  Work-around: Incremental InPacketDiscards during call setup do not necessarily indicate a problem. However, we recommend you investigate an excessive count because that might indicate a connection that cannot be successfully brought up for some reason, such as RADIUS denials or improper configuration.

- The inPacketOctetDiscards counter in the output of the **show atm vc atm** *interface* **vcd** command includes both inBytesDropped and inBytesUnknownProtocol statistics. The inBytesUnknownProtocol statistics should be displayed by a separate counter.

  At the major interface level, the inPacketDiscards counter includes both inPacketsDropped and inPacketUnknownProtocol statistics. The inPacketUnknownProtocol statistics should be displayed by a separate counter. [Defect ID 44286]

BFD

- After you have shut down the interface to the next hop (for the route that is used to establish the BFD session), output for the **show bfd session** command erroneously indicates the shutdown interface as Management Interface (FastEthernet 6/0). [Defect ID 174271]

CLI

- When you issue a **run show ppp** command, the CLI changes the configuration level of the command line to Global Configuration mode rather than remaining at the level from which you issued the command. [Defect ID 52165]

  Work-around: Reissue the commands necessary to reenter the desired mode.

- The **logout subscribers all** command may not log out all of the DHCP subscribers. Although the bindings and DHCP addresses are cleared, the **show subscribers summary** command may display some of the DHCP subscribers. [Defect ID 180176]

  Work-around: Try using the **dhcp delete-binding all** command. If this does not clear the subscribers, you may want to reload the line module to avoid further issues.

DHCP

- DHCP packets are not forwarded to the DHCP server over dynamically created interfaces when all of the following are true: [Defect ID 180343]
  - DHCP relay or DHCP relay proxy is configured on the router.
  - The client-facing interfaces are created dynamically using bridged Ethernet over static ATM PVCs.
  - The **ip auto-detect ip-subscriber** command is configured to enable packet detection (packet triggering) and to trigger creation of dynamic subscriber interfaces.

  Work-around: To avoid this defect, do all of the following:
  - Do not use the **ip auto-detect ip-subscriber** command to enable packet triggering and to create dynamic subscriber interfaces
  - Ensure that DHCP external server is configured in the virtual router.
  - Ensure that the **set dhcp relay inhibit-access-route-creation** command is configured in the virtual router to prevent DHCP relay from installing host routes by default.

DHCP External Server

- With the unique client ID option enabled, when two clients with the same MAC address or client ID are on an interface (where one client is connected over a router and relay and the other client is connected directly), sending a release request from one of the clients might terminate another client. [Defect ID 179759]
- The DHCP renew counter and release counter (displayed with the **show ip dhcp-external statistics** command) are doubled rather than incremented for each renew and release sent. [Defect ID 78802]
- When DHCP clients on an S-VLAN over bridged Ethernet stack configuration send a decline message to a router that has DHCP relay and DHCP external server configured in the same VR, the clients bindings are not removed from the DHCP external server. [Defect ID 87086]
- When DHCP relay and DHCP external server are configured in the same VR with server-sync enabled, bindings are not created in the DHCP external server when DHCP clients on an ATM bulk configuration interface stack and dynamic VLAN over Ethernet stack sends a renew message. [Defect ID 87087]

DoS Protection

- A Telnet session closes when sending ipLocalBGP protocol traffic at a rate in the range 4096–4200 packets per second (pps) with suspicious control flow detection enabled. [Defect ID 81974]

  Work-around: When the traffic drops below 4096 pps, open a new Telnet session.

Ethernet

- When autonegotiation is enabled on Gigabit Ethernet interfaces with the **speed automatically negotiate** command, issuing the **link selection** command logs out subscribers. [Defect ID 87185]

  Work-around: Use the following commands to enable auto link selection (GE port redundancy) and to switch from one port to the other port:

  host1(config-if)#**no link selection**
  host1(config-if)#**link failover force**

File System

- When the primary SRP module is running JunosE Release 7.2.0 or higher-numbered release and the standby SRP module is running a release numbered lower than Release 7.2.0 (as in a downgrade situation), you cannot display the files for the standby SRP module. [Defect ID 74104]

Forwarding

- When performing MAC validation to match subscriber demux entries with ARP host entries, the ES2 10G LM does an exact match, rather than a longest prefix match. The subscriber demux entry source address must be a /32 value matching the IP address of an ARP entry in order to validate the MAC address against that ARP entry. [Defect ID 79641]

- When PPPoE over LAG is configured on an interface, and you re-execute the PPPoE-over-LAG configuration before you delete the previous configuration, the ES2 10G LM line module resets. [Defect ID 179639]

  Work-around: Before you can re-execute the PPPoE-over-LAG configuration, delete the existing PPPoE-over-LAG configuration.

- Specifying S-VLAN ranges that partially overlap does not work. [Defect ID 81918]

  For example, the following configuration fails because S-VLAN 22 falls within the previously specified S-VLAN range of 21–23.

  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 21 23 401 426**
  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 21 23 427 712**
  host1(config-if)#**vlan bulk-config BulkCezarCnfg2 svlan-range 22 22 101 110**

  Work-around: You can do either of the following to avoid this problem.

  – Specify each S-VLAN within the partially overlapping range as individual S-VLANs, as in the following example:

  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 21 21 401 426**
  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 22 22 401 426**
  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 23 23 401 426**
  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 21 21 427 712**
  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 22 22 427 712**
  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 23 23 427 712**
  host1(config-if)#**vlan bulk-config BulkCezarCnfg2 svlan-range 22 22 101 110**

  – Use fully overlapping ranges rather than partially overlapping ranges, as in the following example:

  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 21 23 401 426**
  host1(config-if)#**vlan bulk-config BulkDHCPCnfg1 svlan-range 21 23 427 712**
  host1(config-if)#**vlan bulk-config BulkCezarCnfg2 svlan-range 21 23 101 110**

- When you attach certain hierarchical policies to subinterfaces as input policies, secondary input policies, and output policies, incoming traffic loss can occur when the number of subinterfaces to which the policies are attached exceeds 4600. [Defect ID 86741]
- Ethernet statistics are incorrectly displayed for virtual port 8 of the ES2-S1 GE-8 IOA when that module is paired with the ES2 10G LM or the ES2 10G Uplink LM. [Defect ID 174784]
- The ES2 10G LM does not support framed routes configured for dynamic subscriber interfaces. [Defect ID 83154]
- On the ES2 10G LM, a VLAN ID of 0 assigned to an interface can prevent packets from being properly forwarded. [Defect ID 176125]

ICR

- If you saved the running configuration of the router as a script file (.scr) and execute the script to apply the settings on the router, ICR partition configuration commands in the .scr file might fail to add group members to the partition. This problem happens when the subscriber configuration in the .scr file is placed before the ICR partition configuration. However, this problem does not occur if you used a system configuration (.cnf) file to set up the router. [Defect ID 183913]

Work-around: To correct this problem and enable ICR partitions to be created correctly, make sure that you add the ICR partition configuration before the subscriber interface configuration in the .scr file. You can perform this reordering by modifying the .scr file to place the commands that configure subinterfaces for ICR partitions before the commands used for VLAN-based or S-VLAN-based grouping of subscribers.

- When you configure ICR settings using a CLI macro, ICR commands are run in quick succession. Sometimes, in such a scenario, the active SRP module resets if the event that causes the change of state of the VRRP instance reaches the ICR application before the ICR partition has been created. [Defect ID 184095]

  Work-around: To avoid this problem, add an additional delay of one second using the **sleep** command in the macro, before the **ip vrrp vrid enable** command that is written in the macro to enable VRRP instance.

  For example, consider a macro that contains the following commands:

  ip vrrp *vrid* **enable**
  ip vrrp *vrid* **icr-partition** *partitionId*

  Modify the macro, as follows, to add a delay of one second before the VRRP instance ID is enabled on the router and a delay of another second before the ICR partition that corresponds to the VRRP instance is created:

  sleep 1
  ip vrrp *vrid* **enable**
  sleep 1
  ip vrrp *vrid* **icr-partition** *partitionId*

IGMP

- IGMPv3 proxy is not supported. [Defect ID 46038]
- The default value for the IGMPv3 proxy unsolicited report interval timer should be 1 second rather than 10 seconds (the value for v2). [Defect ID 46040]
- The E Series router IGMPv3 proxy does not operate correctly in the presence of IGMPv2 queriers. [Defect ID 46039/46045]

  Work-around: If an IGMPv2 router is present on the network, do not configure version 3 with the **ip igmp-proxy version** command on that network interface. (Version 2 is the default.)
- When more than about 100,000 mapped OIF entries are configured on a virtual router, issuing the **no virtual router** command for this and other virtual routers does not delete all the virtual routers within the deletion timeout interval (3 minutes). The virtual routers do eventually delete after this timeout. [Defect ID 63882]
- The E Series router does not log a warning when it receives an IGMPv2 query but is not configured to use IGMPv2 on the interface. [Defect ID 46046]

IP

- The ES2 4G LM can reset during a unified ISSU after you issue the **issu start** command on a router configured with 8000 dynamic VCs and 8000 packet-triggered dynamic subscriber interfaces. [Defect ID 86761]

- If you have a large configuration on a hybrid module combination (OC3/STM-1 GE/FE line module with the OC3-2 GE APS I/O module), boot from NVS, and issue the **slot erase** command before booting has completed, the line module resets. [Defect ID 64104]

  Work-around: To recover from the error, issue the **slot reload** command anytime after the module begins to reset.

- Deleting a VRF with 32,000 static subscriber interfaces fails to complete. [Defect ID 82670]

  Work-around: Use a macro to delete all static subscriber interfaces before you delete a VRF.

- The **ip route permanent** command does not work properly. [Defect ID 34303]

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

- Traffic statistics for dynamic subscriber interfaces associated with Mobile IP subscribers are not maintained as the subscribers move between Mobile IP nodes. Consequently the reported interface statistics are only the values accumulated since the last time a mobile node moved. [Defect ID 174509]

- When a router configured with PIM on a virtual router undergoes multiple warm restarts, the router subsequently hangs when an IP profile is configured. [Defect ID 176470]

- Logical port 20 on the ES2-S3 GE-20 IOA is reserved for the hardware multicast packet replication feature. Logical port 20 and the hardware multicast replication feature are not supported on the ES2-S3 GE-20 IOA in this release. [Defect ID 84727]

- When you change the demultiplexer type on a primary interface that has 1024 demultiplexer table entries, the ICC ping threshold times out due to the removal of the old entries and the addition of the new ones. [Defect ID 182218]

- After an SRP stateful switchover completes on an ERX1410 router configured with a single VPN routing and forwarding instance (VRF) and Network Address Translation (NAT), the SRP module that becomes active after the switchover resets. [Defect ID 180058]

IPSec

- In a network where you use the **tunnel signalling** command to specify that the security parameters and keys are configured manually for IPSec tunnels between VRs, the line modules reset when you delete and then re-create the IPSec tunnels. If you attempt to configure the tunnels again after the modules come back up, the line modules reset again.

  Work-around: Configure the IPSec tunnels to use ISASKMP/IKE to negotiate SA and establish keys. [Defect ID 178304]

- IPSec tunnels created over Fast Ethernet interfaces fail to come up. [Defect ID 179256]

  Work-around: After you create the tunnel, bounce the tunnel interface by issuing the **shutdown**/**no shutdown** command sequence. The tunnel comes up successfully.

IS-IS

- On a router configured with IS-IS and BFD, using the **redundancy force srp** command to force an SRP switchover sometimes brings down IS-IS and BFD. [Defect ID 179287]
- IS-IS graceful restart (nonstop forwarding) does not work on the broadcast interface when the restarting router is the designated intermediate system (DIS). Graceful restart works properly when the restarting router is not the DIS. [Defect ID 61496]

L2TP

- After a unified ISSU completes on a router functioning as an L2TP access concentrator (LAC), traffic outages occur on the L2TP network server (LNS)-facing interface at the LAC in a configuration with 16,000 or 32,000 L2TP sessions over 500 tunnels. [Defect ID 180147]

MLD

- MLDv2 proxy is not supported. [Defect ID 46038]
- The E Series router MLDv2 proxy does not operate correctly in the presence of MLDv1 queriers. [Defect ID 46039/46045]

  Work-around: If an MLDv1 router is present on the network, configure version 1 with the **ipv6 mld-proxy version** command on that network interface. (Version 2 is the default.)
- The default value for the MLDv2 proxy unsolicited report interval timer should be 1 second rather than 10 seconds (the value for v1). [Defect ID 46040]
- The E Series router does not log a warning when it receives an MLDv1 query but is not configured to use MLDv1 on the interface. [Defect ID 46046]

MLPPP

- Failure to meet all of the following conditions for fragmented packets can result in an incorrect operation during packet classification of the resulting reassembled packet: [Defect ID 50111]
  - The initial fragment of a packet must either contain the entire MLPPP packet or be greater than 128 bytes.
  - The fragment size of the peer must not be lower than 128 bytes.
  - The initial fragment of a packet must be larger than subsequent fragments of that packet.

Mobile IP

- The **clear ip mobile binding nai @realm** command does not work. [Defect ID 178652]

  Work-around: Use the following version of the command instead:

```
clear ip mobile binding nai user@realm
```

- The *@realm* variable and the **@** keyword alone do not work for the **show ip mobile binding** command. [Defect ID 178653]

  Work-around: You can use the *user@realm* syntax instead to display the binding for a specific user, as in this example:

```
host1#show ip mobile binding nai xyz@example.com
```

Alternatively, you can display the entire Mobile IP binding table by issuing the **show ip mobile binding** command without additional options.

- The setup rate for Mobile IP client sessions decreases when you repeatedly bring a large number of sessions down and back up. [Defect ID 178760]
- When mobility bindings are present and you delete the Mobile IP home agent with the **no virtual router** command, Mobile IP sends a RADIUS Acct-Stop message with no accounting statistics for the subscribers. [Defect ID 179081]

  Work-around: Issue the **clear ip mobile binding all** command before you issue the **no virtual router** command. The **clear** command clears all the MIP subscribers and sends a RADIUS Acct-Stop message with the appropriate accounting statistics for the subscribers.

MPLS

- When MPLS and IS-IS are configured on Ethernet interfaces, you cannot delete the interface after the IP address is removed. This issue is not a problem on Ethernet VLAN interfaces. [Defect ID 66813]

  Work-around: Issue the **no mpls** command to disable MPLS, then delete the interface.
- If LSPs are announced into IS-IS, then the IS-IS routes cannot be used for multicast RPF checks, because LSPs are unidirectional. [Defect ID 28526]

  Work-around: Configure static RPF routes with native hops when LSPs are autoroute announced to IGPs.
- When the IPv4 explicit null label appears anywhere other than at the bottom of the label stack, TTL expiration for this label is not handled correctly. As a result, the **traceroute** command does not work correctly for LSPs that have the IPv4 explicit null label anywhere other than at the bottom of the label stack. [Defect ID 76037]
- When you issue a **traceroute** or **trace mpls** command to trace the paths of router packets over MPLS interfaces on an ES2 10G LM or ES2 10G Uplink LM, the results include an extra unknown host. [Defect ID 174537]

Multicast

- When you configure more than 10,219 outgoing interfaces (OIFs) on the same ES2 10G LM in a single multicast group, the configuration of the multicast group's OIF membership from the SRP module to the line module exceeds the size of a single message and is sent in fragments. Because of this fragmentation, the ES2 10G LM generates the following error message: [Defect ID 81768]

  pc: 0x9e5c88: -> fatalPanic(void) offset: 0x8

Netflow

- Flow sampling stops after a cold switchover on a router that is configured with 16 VRs and 32 interfaces per VR, when all flows are passing through the configuration (32 flows per VR). [Defect ID 74477]

  Work-around: After the cold switchover is completed, reissue the **ip flow-sampling-mode packet-interval 10** command on each VR, even though the command is present in the configuration.

- The OC3/STM1 GE/FE line module might reset after sending Ethernet traffic into a VPLS network in a test environment when Ethernet packets are flooded to remote VPLS bridges. [Defect ID 74540]

Policy Management

- On the E320 router, redirecting a large configuration with thousands of interfaces to a script file can take a long time, perhaps exceeding a half-hour depending on the configuration. [Defect ID 80429]
- If you have removed the last rule in a policy list, the router generates a warning only after you exit Policy List Configuration mode. If you have removed the last policy rule and then added a classifier group before you exit Policy List Configuration mode, the router does not generate a warning about removing the last rule. [Defect ID 83834]
- When an MD-Port-Number value greater than 65,535 is sent to an E120 or E320 router by means of a COA request, the value that is displayed in the UDP header of mirrored packets is the actual value minus 65,536. For example, an MD-Port-Number of 65,540 is displayed in the mirrored packet as 4. [Defect ID 84712]
- On the E120 and E320 routers, when a mirror rule is deleted after a CoA request is sent with Juniper-LI-Action set to No-Action, the existing mirroring session is not disabled. [Defect ID 84826]
- When you reload the slot holding a GE-2 or GE-HDE line module and you have configured more than about 2000 policies with rate limiting on that module, the drop count becomes more than expected. This unexpected drop count does not occur when you create the same configuration after you reload the router to the factory-default configuration. [Defect ID 175696]
- On E320 line modules that support secure policies, the SRP module enables you to configure more than 1022 secure policies per module. [Defect ID 175756]

  Work-around: To avoid potential performance issues, we recommend that you do not configure more than 1022 secure policies per module.
- Unified in-service software upgrade (unified ISSU) is not supported on an E120 or E320 router if a hierarchical policy is attached to an external parent group. [Defect ID 177478]
- When you modify a rate-limit profile in Global Configuration mode after the system is in a scaled state, changes to the rate-limit profile fail owing to lack of adequate policy resources. However, the changed value of the rate-limit profile is displayed in the output of the **show rate-limit profile** command. [Defect ID 79342]

  Work-around: To avoid this problem, do not update the rate-limit profile in Global Configuration mode in a scaled environment.
- When you enter the **no ip policy-parameter hierarchical parameterName** command or **no ipv6 policy-parameter hierarchical parameterName** command for a hierarchical policy-parameter type in Interface Configuration mode, the explicit reference of the parameter is removed successfully from the interface. However, the Referenced by interfaces field in the output of the **show policy-parameter** command does not change from the previously configured value to implicit. [Defect ID 183957]

  Work-around: To correct this problem, remove the entire interface configuration.

PPPoE

- The E Series router erroneously accepts a PADI with a payload length of 0 instead of rejecting it and incrementing the PPPoE Invalid PAD packet length counter. [Defect ID 48356]

QoS

- You cannot paste a **load-rebalance** command string that uses the percent option into a console or Telnet session from **show configuration** output because the output displays the % sign rather than the **percent** keyword that was submitted with the command and the percent sign is not recognized by the CLI. [Defect ID 81705]
- The compound shared shaping feature does not work properly on egress forwarding ASIC 2 (EFA2)-based ATM line modules when the shared shaper is queue-controlled as opposed to node-controlled. In a node-controlled configuration, in which you configure the shared-shaping rate on the best-effort scheduler node for the logical interface, integration of the EFA2 and ATM segmentation and reassembly (SAR) schedulers functions properly. However, in a queue-controlled configuration, in which you configure the shared-shaping rate on the best-effort queue for the logical interface, integration of the EFA2 and ATM SAR schedulers does not function properly. [Defect ID 69167]

  Work-around: Use node-controlled compound shared shaping configured on the best-effort scheduler node with EFA2-based ATM line modules.
- The CLI erroneously enables you to configure a QoS profile with the **ethernet node group** command. [Defect ID 80861]
- The dynamic shaping rate calculated by the simple shared shaper can vary because of the variation in the enqueue rate of the constituent queues. Even when the offered load is constant, the mechanism that calculates the enqueue rate introduces a slight variation, introducing a slight variation in the calculated dynamic shaping rate. [Defect ID 80938]
- On a router that has both an ES2 10G LM and an ES2 4G LM installed, the byte count reported by the **show fabric-queue egress-slot** command is incorrect. The reported packet count is correct. [Defect ID 80965]
- PPP sessions may be dropped if you change the shaping rate in a QoS profile that affects thousands of circuits while QoS traffic affected by the profile is being forwarded. [Defect ID 82950]

  Work-around: Do not change the shaping rate in a QoS profile that affects thousands of circuits while QoS traffic is using the profile.
- When QoS resources such as failure nodes and statistics bins are exhausted because of insufficient memory available on the line module, the failures are properly logged, but additional log messages are generated every 10 minutes that report zero failures. [Defect ID 85105]
- The **no qos-parameter-define definition** command does not delete the specified QoS parameter definition. [Defect ID 176844]

  Work-around: Remove the interface and add the desired QoS parameters when you re-create the interface instead of deleting the definition.

- When 32,000 subscribers with 128,000 QoS queues are brought up on an ES2 10G or ES2 10G ADV LM, the LM resets if you modify the QoS profile that contains the best-effort IP or VLAN node rule, which references a scheduler profile configured with shared shaping rate, to a scheduler profile configured with legacy shaping rate. [Defect ID 183291]

  Work-around: To avoid this problem, apply shared shaping on the best-effort queue, instead of on the best-effort node.

- Simple shared shaping does not function correctly when it is used for 32,000 subscribers on an ES2 10G ADV LM. However, when you change the shaper to compound shared shaping, it works properly. Also, simple shared shaping does not function correctly for 16,000 subscribers on an ES2 10G ADV LM. [Defect ID 183512]
- When you configure an E120 or E320 router with an ES2 10G ADV LM as a LAC on one side of an L2TP tunnel and as a LNS to receive packets from the LAC on the other side of the tunnel, use RADIUS servers for authentication of subscribers on both sides of the tunnel, and attempt to bring up 16,000 subscribers on the L2TP tunnel, the LM that has subscribers on the LAC side of the tunnel resets when approximately 8000 logged-in subscribers are logged out and try to reestablish the connection. [Defect ID 184118]

RSVP-TE

- After stateful SRP switchover, forwarding of VPN traffic might not resume if the core interface that carries an MPLS base tunnel with LDP over RSVP-TE flaps (constantly goes up and down). [Defect ID 182019]

Server Card Manager (SCM)

- High availability mode transitions to the pending state when you perform the following steps. The high availability state of the system is displayed in the output of the **show redundancy detail** command.

  1. Configure a shared tunnel-server port on an ES2 4G line module that functions as the primary in a redundancy group of line modules.
  2. Bring up a GRE tunnel on the primary line module.
  3. Perform a line module redundancy operation to switch over from the currently active primary to the standby module.

  When the system is in the pending state, the SCM application running on the router becomes unsupported for 5 minutes, and then it returns to the active state. The client field in the output of the **show redundancy clients** command displays the status of the SCM application. [Defect ID 188489]

Service Manager

- After you activate an independent IPv6 service and issue either of the following commands on the default virtual router or any other virtual router, except the one on which the subscriber session is active, no output is displayed in the CLI interface: [Defect ID 181929]

  – **show service-management subscriber-session** *subscriberName* **interface** *interfaceType interfaceSpecifier*

  – **show service-management subscriber-session** *subscriberName* **interface** *interfaceType interfaceSpecifier* **service-session** *serviceName*

This problem also occurs when a subscriber is authenticated using a RADIUS server for a combined IPv4 and IPv6 service in a dual stack.

Work-around: To avoid this problem, use the **show service-management owner-session** *ownerName ownerId* command to display subscriber session information based on the session owner, instead of the **show service-management subscriber-session** *subscriberName* **interface** *interfaceType* command to display details on subscriber sessions.

SONET

- You cannot use the highest sensitivity bit-error rate setting (a value of 9) associated with APS/MSP alarm when you issue the **threshold sd-ber** command to configure a cOCx/STMx line module with cOC12-APS-capable IOAs. [Defect ID 72861]

  Work-around: Use only a value in the range 5–8 when you issue the **threshold sd-ber** command for this module combination, as in the following example:

  host1(config)#**controller sonet 2/1**
  host1(config-controll)#**aps group boston**
  host1(config-controll)#**aps protect**
  host1(config-controller)#**threshold sd-ber 6**

SRC Software and SDX Software

- When multiple IPv6 interfaces are configured with policies attached from SRC, only some of the IPv6 interfaces have the policies attached. [Defect ID 179498]
- Changing the SSCC status (enable/disable) while IPv6 interfaces are configured might cause the SRP to reset. [Defect ID 179537]

Stateful SRP Switchover (High Availability) and IP Tunnels

- A packet loss sometimes occurs during stateful SRP switchover when you use the **ping** command on a router that is configured for OSPF graceful restart, and is connected to a helper router in the OSPF IPv6 broadcast network and another helper router in the OSPF IPv6 backbone area. [Defect ID 181470]
  - ERX7xx model, ERX14xx model, or ERX310 router:
    - When you use the **ping** command with the IPv6 address of the helper router in the multicast area as the destination address and the loopback address of the helper router in the backbone area as the source address, a packet loss of 2 seconds occurs for the first stateful SRP switchover. However, no packet loss occurs for successive stateful SRP switchovers.
    - When you use the **ping** command with the IPv6 address of the helper router in the broadcast network as the destination address and no source address when stateful SRP switchover is performed the first time, an identical packet loss occurs. In this case too, no packet loss occurs during subsequent switchovers.

  - E120 router or E320 router
    - When you use the **ping** command with the IPv6 address of the helper router in the broadcast network as the destination address and the loopback address of the helper router in the backbone area as the source address, no packet loss occurs.
    - When you use the **ping** command with the IPv6 address of the helper router in the multicast area as the destination address and no source address, a packet loss of 1–2 seconds sometimes occurs during stateful SRP switchovers.
- On E120 and E320 routers configured with an SRP module that contain a high availability pair of line modules, the primary SRP module intermittently resets when you perform a stateful SRP switchover after a stateful line module switchover is completed. This problem occurs only when login and logout of subscribers is in progress during the stateful line module switchover. [Defect ID 186836]

Subscriber Management

- When a dynamic GRE tunnel interface for Mobile IP relocates between SM modules because the original SM reloads, Mobile IP deletes the relocated tunnel interface. [Defect ID 178399]
- When a subscriber has subscribed for a service, service session accounting records always contains a default Acct-Terminate-Cause value of 10. This value remains unchanged even after you use the **terminate-code** command to configure a custom mapping between application terminate reasons and RADIUS Acct-Terminate-Cause attributes. [Defect ID 181043]
- Dynamic subscriber interfaces continue to remain in the down or not present operational state in either of the following scenarios: [Defect ID 81269]
  - If you configured a dynamic interface column, such as a dynamic bridged Ethernet interface, dynamic VLAN interface, or an ATM interface, and when any one of the following conditions is satisfied:
    - The major interface is bounced (shut down and reenabled)
    - The major interface is shut down, which cause the dynamic VLAN interfaces to be removed
    - The physical link goes down and comes back up
    - The line module is removed and reinserted
  - If you configured a static interface column and removed the major interface

  These scenarios might occur if you administratively issue the **shutdown** and **no shutdown** commands on the major interface in which the dynamic interface column is configured.

  Work-around: Use the **no interface ip ipAddress** command to remove the dynamic subscriber interfaces. Although you can use the **dhcp delete-binding** command to remove the DHCP binding and the dynamic subscriber interfaces, the DHCP client does not detect the binding removal and retains the lease.

System

- You cannot use a configuration script to boot the E320 router. [Defect ID 80304]

- If you hot swap an IOA and then remove it again before that IOA's OK or FAIL LED is illuminated, the associated line module can reset. [Defect ID 177313/177267]

  Work-around: Ensure that you firmly insert the IOA into the chassis when you hot swap IOAs. Do not attempt a second hot swap of an IOA that has not indicated that it completed the first hot swap cycle. You can remove the IOA when either its OK or FAIL LED is illuminated.

- If your router is in Manual Commit mode, then you must issue the **write memory** command before you perform an SRP module switch or a manual reload. You must do this even when you have made no changes to the system configuration and the file systems are synchronized. [Defect ID 44469]

TCP

- The SRP module resets in any of the following circumstances on an E320 router that has a line module configured with 5000 ANCP adjacencies: [Defect ID 176916]
  - When you issue the **issu initialization** command from the console and then reload the line module from a Telnet session.
  - When the client that has the 5000 ANCP clients resets or an intermediate switch resets.
  - When you reload the line module.

Unified ISSU

- ATM line modules might reset after a unified ISSU when you attempt to add memory to a VLAN subinterface in a large bridged Ethernet configuration. [Defect ID 178798]
- Under certain conditions, a unified ISSU from JunosE Release 9.2.0p1-0 to the current release fails, and causes the SRP module and the ES2 4G LM to reset. [Defect ID 179975]
- When any of the subsystems is excluded for a JunosE release, a unified ISSU to that release fails to apply conversion code to all of the line modules. As a result, the line modules reset when they come up with that release. [Defect ID 179595]

  Work-around: To prevent the exclusion of a subsystem file from the release, do the following before you upgrade to a new JunosE release that supports unified ISSU:

  1. Issue the **show subsystems file** *fileName.rel* command, where *fileName* is the name of the software release file, to determine whether any of the subsystem files are excluded from the release.
  2. For each subsystem file that is excluded, issue the **no exclude-subsystem** *subsystemName* command to remove the exclusion for the specified subsystem file.

  If you copied the software release to the router before removing the subsystem file from the exclusion list, you must copy the release to the router again to ensure that all subsystem files are included in the release.

- Unified ISSU is not supported with 8000 bridged Ethernet interfaces on an OC3/STM1 GE/FE ATM line module. [Defect ID 178811/178797/179547]

- During the unified ISSU operation, if you modify the router configuration after the initialization phase of the process is completed and before you issue the **issu start** command to commence the upgrade phase of the unified ISSU process, the unified ISSU procedure completes successfully and the stateful SRP switchover process begins to synchronize between the active and standby SRP modules. When the synchronization process is in progress, the standby SRP module reloads for the second time. After the second reload of the standby SRP module ends, the synchronization process also ends properly.

  Although the standby SRP module reloads for the second time when it is synchronized with the upgraded release, normal router operations, such as handling of subscriber sessions and forwarding of traffic, remain unaffected. [Defect ID 185517]

Resolved Known Problems

The following problems were reported open in Release 11.1.1 and have been resolved in this release, or have been resolved since the 11.1.1 FRS release. For more information about particular resolved problems, you can log in to the JunosE Knowledge Base at https://www2.juniper.net/kb/, enter the defect ID number in the Search by Keyword field, and click Search.

ATM

- PPPoEoA packets 1492 bytes cause ATM re-assembly errors in client side in the 10-2-0p0-2 release. [Defect ID 90782]
- Performance issues after upgrading ATM cards to EFA2 architecture. [Defect ID 90861]

DHCP Client

- UnACKnowledged DHCP RENEW request in DHCP Proxy Client leaks FSM pool memory. [Defect ID 90891]

DHCP Server

- DHCPv6-LS must discard Release messages sent on unicast dest address and send reply to comply with RFC. [Defect ID 90913]

Ethernet

- VLAN Interface Monitoring reports incorrect speed and/or accepts invalid interface name as a parameter. [Defect ID 89749]
- After the primary lag link is shutdown, the dynamic VLAN is deleted instead of being made dormant when subscribers log out. [Defect ID 90815]
- "stacked vlan anyUntagged" with Bridge-ethernet configuration gets accepted in "interface configuration" [Defect ID 91071]
- Link LED of GE/FE line card does not go off if link negotiation both duplex full is configured. [Defect ID 91002]

Forwarding

- LM4 Packets are not getting marked correctly on SVLAN 802.1p bits on E320. [Defect ID 90962]

Hardware

- LM4 reset type: processor exception 0x100 (multiple=0x2ff) task: interrupt [Defect ID 90219]
- New routes cannot be downloaded to LM10 due to tpub problem. [Defect ID 90291]
- (DT) Line card redundancy is broken during HW error and "Not Responding" states at UPGRADE / COLD START / COLD SWITCH. [Defect ID 90836]
- LM10 stops forwarding any traffic. [Defect ID 90207]

HTTP Server

- Malformed HTTP POST request received by HTTP server might cause software panic. [Defect ID 90834]

IP

- Few PPP sessions drop due to "**aaa duplicate check**" failure even though the /32 route is not present in SRP routing table. [Defect ID 90806]

IPv6

- In case of Active Solicitation JunosE does not enter the probe state but keeps toggling between reach and stale. [Defect ID 90801]

ISSU

- E320 SRP reset type: panic arg (0x6) file: osHeap.cc line: 1867 task: configMonitor last errno: 0x380003 [Defect ID 90995]

L2C

- SRP crashed with reset type: exc reboot task "l2c" while running l2cOamTest.mac [Defect ID 184333]

L2TP

- New active SRP reset type halt after HA switch while testing the fix for 90586. [Defect ID 90868]
- L/C does not respond to SCCRQ/stopCCN during HA switch. [Defect ID 90875]
- GE-2: reset type: panic file: ic1L2tpDriver.cc task: scheduler indicateSfPeerlessNfepSession [Defect ID 90729]
- A new L2TP session cannot be established when the tunnel controlChannel.congestionWindow size reaches 32768. [Defect ID 90864]
- NTT: E320:10-2-1p0-1 release has wrong MIB "Default SNMP ifType" of L2TP interface. [Defect ID 90881]

Multicast

- The following commands, which appear in IPv6 PIM Data MDT Configuration mode, have never been supported and have been removed from IPv6 PIM Data MDT Configuration mode:
  - **ipv6 pim data-mdt**
  - **ipv6 pim join-filter**
  - **ipv6 pim query-interval**
  - **mdt-data-delay**
  - **mdt-data-holdown**
  - **mdt-data-timeout**
  - **route map**
  - **show ipv6 pim data-mdt**
  - **tunnel group-address-pool**
  - **tunnel source**

  IPv6 PIM Data MDT Configuration mode has also been removed from the CLI. [Defect ID 83791]

OSPF

- L2TP sessions go down during ISSU. [Defect ID 187195]

OSPFv3

- Unknown ls-type is generated in OSPFv3 LSA when ERX is rebooted. [Defect ID 89946]

PPPoE

- Subscribers fail to connect after ISSU upgrade because ERX does not reply with PADS for PADRs received. [Defect ID 90914]

QoS Manager

- QoS Profiles modification for VP nodes shouldn't be allowed for NMBA interfaces. [Defect ID 90950]

SSC Client

- Memory leak in the pending request table in SSC client during heavy interface flapping. [Defect ID 91053]

System

- LC-HA:SRP crash; file: ectionManager.cc(line: 3173) task:cm [Defect ID 176808]

Errata

This section identifies errors found in the JunosE documentation. These errors are corrected in subsequent releases of the affected documentation.

- The *Resolved Known Problems* section of the *JunosE 10.2.2 Release Notes* failed to mention in the *IP* section that defect ID 90081, High CPU utilization, was fixed in that release.
- The JunosE documentation for Release 6.0.5 and higher-numbered releases states that when you upgrade the JunosE Software from Release 5.1.1 or lower-numbered releases, you must perform the upgrade in two stages: first to an intermediate release and then to the higher-numbered release that you want to run. This statement is only partially correct; you must perform a two-stage upgrade only when you upgrade from a new NVS card. This restriction is not applicable if you upgrade your software remotely through Telnet or FTP.

  The imprecise information appears in the following JunosE documents:

  - The *Upgrading to JunosE Software Release 6.x.x or Higher-Numbered Releases from Release 5.1.1* or *Lower-Numbered Releases* special hardware notice dated 31 March 2006
  - The *Upgrading to JunosE Software Release 6.x.x or Higher-Numbered Releases from Release 5.1.1 or Lower-Numbered Releases* section in the *JunosE Release Notes* for the following releases:
    - Releases 6.0.5, 6.1.4, and 6.1.5
    - Releases 7.1.2, 7.1.3, 7.1.4, 7.2.x, 7.3.x
    - Releases 8.x, 9.x, 10.x, 11.0.x, and 11.1.x
  - The *Upgrading to JunosE Software Release 6.x.x or Higher-Numbered Releases from Release 5.1.1 or Lower-Numbered Releases* section in *ERX Hardware Guide, Chapter 8, Maintaining ERX Routers* for the following releases:
    - Release 7.3.x
    - Releases 8.x, 9.x, 10.x, 11.0.x, and 11.1.x
  - The *Upgrading to JunosE Software Release 6.x.x or Higher-Numbered Releases from Release 5.1.1 or Lower-Numbered Releases* section in the *JunosE System Basics Configuration Guide, Chapter 3, Installing JunosE Software* for the following releases:
    - Release 7.3.x
    - Releases 8.x, 9.x, 10.x, 11.0.x, and 11.1.x
- The *Release Highlights* section of the *JunosE 10.1.0 Release Notes* failed to state that the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router do not support hot-swap of the SRP IOA.
- The Physical and Logical Density Maximums table (for the E120 and E320 routers) in *Appendix A, System Maximums* of the *JunosE Release Notes*, for the following releases incorrectly contained a note stating that you could install the ES2 10G ADV LM in either of the E320 router turbo slots (2 and 4) with a 100 Gbps fabric configuration on the E320 router.
  - Release 10.1.x
  - Releases 10.2.0, 10.2.1, and 10.2.2

  - Releases 10.3.0, 10.3.1, and 10.3.2
  - Releases 11.0.0, 11.0.1, and 11.0.2
  - Releases 11.1.0 and 11.1.1

  The ES2 10G ADV LM cannot be installed in any of the slots with a 100 Gbps fabric configuration on an E320 router. You can configure the ES2 10G ADV LM with only the following module combinations:

  - 120 Gbps fabric configuration on the E120 router
  - 320 Gbps fabric configuration on the E320 router

- In the Module Combinations table in *E120 and E320 Hardware Guide, Chapter 4, Installing Modules*, the SFM-100 row incorrectly indicates that this SFM can be paired with the ES2 10G ADV line module. The SFM-100 module is not compatible with an E320 router that contains an ES2 10G ADV LM.

  Also, the SFM-120 and SFM-320 rows erroneously fail to mention that these SFM modules can be paired with the ES2 10G ADV LM. The SFM-120 module can be installed with an ES2 10G ADV LM on an E120 router in slots numbered 0-5 and the SFM-320 module can be installed with an ES2 10G ADV LM on an E320 router in slots numbered 0-5 and 11-16.

- The *Environmental Requirements* and *Cabling Recommendations* sections in *Appendix B, Installation Guidelines and Requirements,* of the *E120 and E320 Hardware Guide* fail to mention that, based on the optical fiber cables used, you may need to increase the physical space provided for the chassis. The strain relief and bend radius requirements of the optical fiber cables may exceed the specified depth of the chassis and the cable-management bracket.

- The *ERX Module Guide* and the *E120 and E320 Module Guide* incorrectly include information about the following end-of-life (EOL) modules, which can no longer be ordered for ERX routers, E120 routers, and E320 routers:

Module Name	Module Type
cOC3/STM1 multimode	ERX
cOC3/STM1 single-mode intermediate reach	ERX
cOC3/STM1 single-mode long reach	ERX
cOC12/STM4multimode without APS/MSP redundancy	ERX
cOC12/STM4 single-mode intermediate reach without APS/MSP redundancy	ERX
cOC12/STM4 single-mode intermediate reach with APS/MSP redundancy	ERX
cOC12/STM4 single-mode long reach	ERX
CT3/T3 12 (12 ports)	ERX
E3 Frame (12 ports)	ERX
IPSec Service Module	ERX
SRP 5G+(2 GB)	ERX
T3 Frame (12 ports)	ERX
10GE IOA	E120 and E320

- *ERX Module Guide, Appendix A, Module Protocol Support* incorrectly states that IPv6 multicast is supported on the following line modules:
  - cOCx FO line module with cOC3/STM1 modules
  - cOCx FO line module with cOC12/STM4 FO I/O modules
  - COCx-F3 line modules with CT3/T3 12 I/O modules
  - OCx/STMx ATM line modules with 4xDS3 ATM I/O modules
  - CT3/T3-FO line modules with CT3/T3 12 I/O modules
- The following corrections and additions apply to the LM 10 ADV Line Module table in the *E120 and E320 Module Guide*:
  - The Description row fails to mention that you can use the ES2 10G ADV LM with an SRP-120 configuration and that you can install it in any slot.
  - The Capability row incorrectly states that for the ES2 10G ADV LM, the 100–Gbps switch fabric allocates 3.4 Gbps of overall bandwidth to each regular LM slot and 10 Gbps of overall bandwidth to each of the turbo slots (slot 2 and slot 4). However, the correct compatibility setting is that the ES2 10G ADV LM cannot be paired with a 100–Gbps switch fabric. You can use this LM only with SRP-120 and SFM-120 modules on an E120 router and with SRP-320 and SFM-320 modules on an E320 router.

    Also, the Capability row fails to mention that when you use the ES2 10G ADV LM with an SFM-120 module, the 120–Gbps switch fabric allocates 10 Gbps of overall bandwidth to each LM slot.
  - The SRP module compatibility row fails to mention that you can use the ES2 10G ADV LM with an SRP-120 module.
- The Line module compatibility row in the SRP 100 Module and SFM 100 Module tables of the *E120 and E320 Module Guide* incorrectly mentions that the ES2 10G ADV LM (only in turbo slots) can be installed with the SRP-100 and SFM-100 modules. The ES2 10G ADV LM is not compatible with both the SRP-100 and SFM-100 modules.
- The Line module compatibility row in the SRP 120 Module and SFM 120 Module tables of the *E120 and E320 Module Guide* fails to list the ES2 10G ADV LM as one of the LMs that can interoperate and be paired with the SRP-120 and SFM-120 modules.
- *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* incorrectly states that IPv6 multicast is supported on the ES2-S1 Service IOA module. IPv6 multicast is not supported on this module.
- The *ERX7xx Models, ERX14xx Models, and the ERX310 Broadband Services Router* section of *JunosE System Basics Configuration Guide, Chapter 6, Managing Modules* failed to mention that the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router do not support hot-swap of the SRP IOA.
- *The show clns interface* section in *JunosE IP, IPv6, and IGP Configuration Guide, Chapter 6, Configuring IS-IS*, incorrectly indicates that LDP-IGP Synchronization is supported only for OSPFv2. In fact, LDP-IGP Synchronization is supported for IS-IS.

- The *Detecting Corrupt File Configurations* section in *JunosE System Basics Configuration Guide, Chapter 5, Managing the System* fails to mention that if you check the running configuration for corruption manually when auto mode is enabled, a warning message appears:

```
host1(config)#service check-config running-configuration
WARNING: This command will cause config monitor to switch into manual mode.
Proceed with current command? [confirm]
```

  If you confirm you want to check the running configuration in manual mode or press any key to continue, then manual mode is enabled.

- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 4861—Neighbor Discovery for IP Version 6 (IPv6) (September 2007). RFC 4861 obsoletes RFC 2461—Neighbor Discovery for IP Version 6 (IPv6) (December 1998). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.
- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 4862—IPv6 Stateless Address Autoconfiguration (September 2007). RFC 4862 obsoletes RFC 2462—IPv6 Stateless Address Autoconfiguration (December 1998). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.
- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 4760—Multiprotocol Extensions for BGP-4 (January 2007). RFC 4760 obsoletes RFC 2858—Multiprotocol Extensions for BGP-4 (June 2000). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.
- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 4007—IPv6 Scoped Address Architecture (March 2005). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.
- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 4193—Unique Local IPv6 Unicast Addresses (October 2005). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.
- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 2526—Reserved IPv6 Subnet Anycast Addresses (March 1999). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.
- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 5072—IP Version 6 over PPP (September 2007). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.
- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 5187—OSPFv3 Graceful Restart (June 2008). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.
- *JunosE System Basics Configuration Guide, Appendix B, References*, fails to list RFC 5308—Routing IPv6 with IS-IS (October 2008). Since the publication of the RFC, the JunosE Software has been compliant with the updated RFC.

- In *JunosE IP, IPv6, and IGP Configuration Guide, Chapter 2, Configuring IPv6*, the *Before You Configure IPv6* section fails to list all the modules that support IPv6. You can find complete lists of modules that support IPv6 in the following appendixes:
  - *E120 and E320 Module Guide, Appendix A, IOA Protocol Support*
  - *ERX Module Guide, Appendix A, Module Protocol Support*
- In the *Configuring Graceful Restart* section in *JunosE IP, IPv6, and IGP Configuration Guide, Chapter 6, Configuring IS-IS*, the default time for the restarting router to wait for the LSP database to synchronize and to reset the overload bit is incorrectly mentioned as 30 seconds. The correct value is 100 seconds for both the instances.
- The *Monitoring SONET/SDH Interfaces* section in *JunosE Physical Layer Configuration Guide*, *Chapter 3, Configuring Unchannelized OCx/STMx Interfaces*, and the *Monitoring Interfaces* section in *Chapter 4, Configuring Channelized OCx/STMx Interfaces*, incorrectly display the field description of the "time since last status change" field as "time since the module was rebooted" for the **show controllers sonet** and **show controllers sonet configuration** commands.

  The correct field description is "time the controller has been in the current physical state" for all four instances.
- The *ES2 10G ADV LM Combination* section in *JunosE Physical Layer Configuration Guide, Chapter 5, Configuring Ethernet Interfaces* incorrectly states the following:

  With a 100 Gbps fabric configuration, the E320 router can accommodate up to 2 combinations of ES2 10G ADV LMs and ES2-S3 GE-20 IOAs. You must install a combination in either of the turbo slots (slot 2 or slot 4). The 100 Gbps allocates 10 Gbps of overall bandwidth to each of these slots.

  The correct compatibility configuration is that you cannot install an ES2 10G ADV LM in any of the slots with a 100 Gbps fabric configuration on an E320 router. You can configure the ES2 10G ADV LM with only the following module combinations:
  - 120 Gbps fabric configuration on the E120 router
  - 320 Gbps fabric configuration on the E120 and E320 routers
- The *Configuring a Profile* section of *JunosE Link Layer Configuration Guide, Chapter 17, Configuring Dynamic Interfaces*, incorrectly mentions the **pppoe motm** command as the **pppoe mtu** command.

  The command name is now replaced with the **pppoe motm** command. Also, the description of the **pppoe motm** command is updated with additional information.
- In the *Configuring a Profile* section in *JunosE Link Layer Configuration Guide, Chapter 17, Configuring Dynamic Interfaces*, the **ppp authentication** command section erroneously omits the following information from the list of guidelines:

  If the VR specified in a profile with the **ipv6 virtual-router** command differs from the VR provided by AAA, IPv6 uses the VR provided by AAA when the dynamic IPv6 upper-layer interface is created.
- In the *JunosE Service Availability Configuration Guide* for Releases 10.3.0, 11.0.0, and 11.1.0, the Notes column for the Multicast Routing rows under IPv4 Multicast Routing and under IPv6 Multicast Routing in the *Application Support for Stateful SRP Switchover* table in *Chapter 3, Managing Stateful SRP Switchover*, stated incorrectly that only static recovery is supported. The correct information is as follows:

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

Informational Note: Before JunosE Release 10.3.0, the Application Support for Stateful SRP Switchover table was located in the *JunosE System Basics Configuration Guide*.

- The *Grouping ICR Subscribers Based on S-VLAN IDs* and *Example: Configuring ICR Partitions That Group Subscribers by S-VLAN ID* sections in *JunosE 11.1.x Service Availability Configuration Guide, Chapter 6, Managing Interchassis Redundancy* fail to mention that, to enable the master router to send PPPoE Active Discovery Termination (PADT) packets to the clients and create new sessions for the PPPoE subscribers, you must create a dummy IP interface for each S-VLAN subinterface that is part of the ICR partition.
- In the *Grouping ICR Subscribers Based on S-VLAN IDs* and *Example: Configuring ICR Partitions That Group Subscribers by S-VLAN ID sections in JunosE 11.1.x Service Availability Configuration Guide, Chapter 6, Managing Interchassis Redundancy*, the **use-default-mac** keyword, which is used with the **svlan-list**, **svlan-range**, and **svlan-list-explicit** commands, was incorrectly documented as required for generating Gratuitous ARP (GARP). In fact, the keyword **use-default-mac** is not required for generating GARP. An example of the correct usage of the commands is as follows:

  host1(config-if)#ip vrrp 1 icr-partition svlan-list 100 102 105 108 control-interface advertise-mac

  host1(config-if)#ip vrrp 1 icr-partition svlan-range 100 110 control-interface advertise-mac

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

- In the *Monitoring IGMP* section in *JunosE Multicast Routing Configuration Guide, Chapter 2, Configuring IGMP*, the bulleted list of field descriptions and output example for the **show ip igmp interface** command incorrectly display the field name for the number of multicast groups that the interface has discovered as "Groups learned". The correct label for this field is "Groups learnt".

- In the *Creating Multicast VPNs Using Default MDT* section in *Chapter 3, Configuring PIM for IPv4 Multicast* of the *JunosE Multicast Routing Configuration Guide*, the command line configuration examples in Step 8, Step 9, and Step 10 are incorrect. The following steps present the correct information:
  - Step 8: Configure the IP interface (Tv) in PE2:CE1 as a numbered or unnumbered PIM sparse-mode interface. Use the same address as the loopback interface, Lp in the parent router, PE2.

```
host1(config)#virtual-router PE2:CE1
host1:PE2:CE1(config)#interface tunnel gre:MTI-21
host1:PE2:CE1(config-if)#ip address 2.2.2.2 255.255.255.255
host1:PE2:CE1(config-if)#ip pim sparse-mode
host1:PE2:CE1(config-if)#exit
host1:PE2:CE1#
```

  - Step 9: Configure the IP interface (Tv) in PE2:CE2 as a numbered or unnumbered PIM sparse-mode interface. Use the same address as the loopback interface, Lp in the parent router, PE2.

```
host1(config)#virtual-router PE2:CE2
host1:PE2:CE2(config)#interface loopback 0
host1:PE2:CE2(config-if)#ip address 2.2.2.2 255.255.255.255
host1:PE2:CE2(config-if)#exit
host1:PE2:CE2(config)#
host1:PE2:CE2(config)#interface tunnel gre:MTI-22
host1:PE2:CE2(config-if)#ip unnumbered loopback 0
host1:PE2:CE2(config-if)#ip pim sparse-mode
host1:PE2:CE2(config-if)#exit
host1:PE2:CE2#
```

  - Step 10: Configure the Tp interfaces in the parent router, PE2, as unnumbered PIM sparse-mode interfaces tied to the loopback interface, Lp.

```
host1(config)#virtual-router PE2
host1: PE2(config)#interface tunnel gre:MTI-21.mdt
host1:PE2(config-if)#ip unnumbered loopback 0
host1:PE2(config-if)#ip pim sparse-mode
host1:PE2(config-if)#exit
host1:PE2(config)#
host1:PE2(config)#interface tunnel gre:MTI-22.mdt
host1:PE2(config-if)#ip unnumbered loopback 0
host1:PE2(config-if)#ip pim sparse-mode
host1:PE2(config-if)#exit
host1:PE2(config)#
```

- In the *Monitoring MLD* section in *JunosE Multicast Routing Configuration Guide, Chapter 6, Configuring Multicast Listener Discovery*, the following corrections apply to the commands used to monitor MLD configuration:
  - In the **show ipv6 mld** command section, the field name "learned groups" (which denotes the number of multicast groups that the virtual router has discovered) in the bulleted list of field descriptions and output example for this command is incorrect. The correct field name for this setting is "learnt groups".
  - In the **show ipv6 mld interface** command section, the field name "Groups learned" (which denotes the number of multicast groups that the interface has discovered) in the bulleted list of field descriptions and output example for this command is incorrect. The correct field name for this setting is "Groups learnt".

- In the *Aggregating Routes* section in *JunosE BGP and MPLS Configuration Guide, Chapter 1, Configuring BGP Routing*, the references to the router, “Snakes”, in the description that precedes the “Configuring Aggregate Addresses” figure are incorrect. The correct references for this router are “SanJose”, which is the label used for this router in the figure.
- In the *Detecting Peer Reachability with BFD* section in *JunosE BGP and MPLS Configuration Guide, Chapter 1, Configuring BGP Routing*, the **neighbor-bfd-liveness-detection** command subsection incorrectly states the following:

  If you remove the BFD configuration while the BGP sessions and the BFD protocol session are up, then the BGP session may flap because the remote BGP speaker cannot detect why the BFD session went down.

  The correct behavior of BGP sessions, when you remove the BFD configuration for the last client tied to a BFD session, is as follows:

  If you remove the BFD configuration while the BGP sessions and the BFD protocol session are up, BFD moves to the Admin Down state and communicates the change to the peer to enable the client protocols to handle this transition in a seamless manner without going down. For the Admin Down state to work, the peer, which receives the Admin Down state notification, must have the capability to distinguish between administratively down state and real link down.

Informational Note: The BFD Admin Down state is used to bring down a BFD session administratively, to protect client applications from BFD configuration removal, license issues, and clearing of BFD sessions.

- In the *BGP Multihoming for VPLS* section in *Chapter 9, VPLS Overview* of the *BGP and MPLS Configuration Guide,* the Selecting the Designated VE Device for a Multihomed Site algorithm is erroneous. Some references to this algorithm are also incorrect. See *Appendix B, Selecting the Designated VE Device for a Multihomed Site* in these *Release Notes* for updated information.
- In the *Example: Aggregate Marking with Oversubscription Rate-Limiting Hierarchical Policy* section in *Chapter 5, Creating Rate-Limit Profiles* of the *JunosE Policy Management Configuration Guide*, the references to the sample ToS values marked for color coded traffic packets in the description, the command examples, and the figure are incorrect. In this section, the sample values that can be configured for ToS marks, using the **green-mark**, **yellow-mark**, and **red-mark** commands, are specified as a string of characters. The correct ToS marks that can be configured for color-coded packets using these commands are of the numerical format (a value in the range 0–255).

  See *JunosE 11.1.1 Release Notes, Appendix B, Applying a ToS Mark to Color-Coded Packets in Aggregate Rate-Limit Hierarchical Policies* for updated information.

- In the *Attaching a QoS Profile to an ATM VP* section in *JunosE Quality of Service Configuration Guide, Chapter 16, Configuring and Attaching QoS Profiles to an Interface*, the following information about the attachment of QoS profiles to ATM interfaces is missing:

  If you attempt to modify the QoS profile attached to an ATM VP that contains nonbroadcast multiaccess (NBMA) or multipoint interfaces from profileA to profileB by using the **atm-vp qos-profile** command for a specific VP on that interface, the command is configured correctly and no error message is displayed in the CLI interface. However, the shaping rate on the interfaces that are part of the ATM VP is not properly updated with the shaping rate specified in profileB. Instead, the multipoint interfaces remain configured with the shaping rate set in profileA.

  To modify the QoS profile currently attached to ATM VPs that contain NBMA or multipoint interfaces from another profile, you must first remove the QoS profile attached to the interfaces by using the **no atm-vp qos-profile** command in Interface Configuration mode, and then attach the new QoS profile to the interfaces by using the **atm-vp qos-profile** command. This restriction exists because the mungeing of QoS profiles does not occur correctly if any of the attributes of ATM VPs with multipoint interfaces are modified.

  If you modify the QoS profile attached to a point-to-point ATM interface from profileA to profileB by using the **qos-profile** command (or the **atm-vp qos-profile** command for a specific VP on the ATM interface) in Interface Configuration mode, the shaping rate is correctly configured on the interface and is modified with the value specified in profileB.

  To modify the QoS profile attached to an ATM VP that contains an NBMA or a multipoint interface from profileA to profileB, perform the following steps. These steps assume that profileA and profileB have been previously configured on the router.

  1. Enter Interface Configuration mode for the ATM VP.

     host1(config)#interface atm 1/0

  2. Remove the QoS profile, profileA, currently attached to the ATM VP that contains the NMBA interface.

     host1(config-if)#no atm-vp 1 qos-profile profileA

  3. Attach the new QoS profile, profileB, that you want to be attached to the ATM VP that contains the NBMA interface.

     host1(config-if)#atm-vp 1 qos-profile profileB

- The *IP Hinting* section in *JunosE Broadband Access Configuration Guide, Chapter 1, Configuring Remote Access* inadvertently omits the following information about AAA domain map support for IPv4 that enables you to configure additional virtual router assignment capabilities for IPv4 subscribers:

  AAA domain map support for IPv4 enables you to provide additional virtual router assignment capabilities for IPv4 subscribers. If you assign a value other than default to a layer 2 virtual router, then the access, IPv4, and IPv6 virtual routers are all assigned the same value, which cannot be changed. If you use RADIUS redirect to assign virtual routers, you can assign access, IPv4, and IPv6 to the redirection target.

Use the **auth-router-name** command in Domain Map Configuration mode to assign an access virtual router. The **no** version restores the default router. An example of the **auth-router-name** command is as follows:

```
host1(config)#aaa domain-map xyz.com
host1(config-domain-map)#auth-router-name accessvr
```

Informational Note: The **ip-router-name** command replaces the **router-name** command, which has been deprecated and may be removed completely in a future release.

Use the **ip-router-name** command in Domain Map Configuration mode to assign an IPv4 virtual router. The **no** version restores the default router. An example of the **ip-router-name** command is as follows:

```
host1(config)#aaa domain-map xyz.com
host1(config-domain-map)#ip-router-name ipv4vr
```

- The following information was erroneously omitted from the *Generating UDP Checksums in Packets to L2TP Peers* section of *Chapter 12, Configuring an L2TP LAC*, of the *JunosE Broadband Access Configuration Guide* for Release 11.1.x.

  L2TP checksum generation support is available on an ES2 10G Uplink LM and an ES2 4G LM only. It is not supported on an ES2 10G LM and an ES2 10G ADV LM. If an ES2 10G LM or an ES2 10G ADV LM is present when L2TP checksum is enabled, the checksum is not calculated and its value is set to zero.

- The syntax of the **no** version of the **ethernet oam lfm link-monitor frame-seconds** command in the *JunosE Command Reference Guide A to M* is incorrect. This erroneous syntax indicates that you are required to specify either the threshold value or the period of time during which error frames are counted (the window).

  The incorrect syntax is as follows:

  no ethernet oam lfm link-monitor frame-seconds { threshold { high { *highFrames* | none } | low *lowFrames* } | window *period* }

  The correct syntax of the **no** version of this command is as follows:

  no ethernet oam lfm link-monitor frame-seconds [threshold { high { *highFrames* | none } | low *lowFrames* } | window *period*]

  Using this syntax, you are not required to specify a threshold value or a time period. Both the **threshold** and **window** keywords are optional when you use the **no** form of this command.

  If you enter the **no ethernet oam lfm link-monitor frame-seconds** command without specifying the threshold or window parameters, both the configured threshold and window attributes are removed. If you want to delete only one of the two settings, such as the threshold value or the period of time for error frames monitoring, you can use the keyword specific to the setting that you want to remove with the **no ethernet oam lfm link-monitor frame-seconds** command.

- The syntax of the **no** version of the **ethernet oam lfm link-monitor frame-seconds-summary** command in the *JunosE Command Reference Guide A to M* is incorrect. This erroneous syntax indicates that you are required to specify either the threshold value or the period of time during which errored frame seconds are counted (the window).

The incorrect syntax is as follows:

no ethernet oam lfm link-monitor frame-seconds-summary { threshold { high { *highFrames* | none } | low *lowFrames* } | window *period* }

The correct syntax of the **no** version of this command is as follows:

no ethernet oam lfm link-monitor frame-seconds-summary [threshold { high { *highFrames* | none } | low *lowFrames* } | window *period*]

Using this syntax, you are not required to specify a threshold value or a time period. Both the **threshold** and **window** keywords are optional when you use the **no** form of this command.

If you enter the **no ethernet oam lfm link-monitor frame-seconds-summary** command without specifying the threshold or window parameters, both the configured threshold and window attributes are removed. If you want to delete only one of the two settings, such as the threshold value or the period of time for monitoring errored frame seconds, you can use the keyword specific to the setting that you want to remove with the **no ethernet oam lfm link-monitor frame-seconds-summary** command.

- The syntax of the **no** version of the **ethernet oam lfm link-monitor symbol-period** command in the *JunosE Command Reference Guide A to M* is incorrect. This erroneous syntax indicates that you are required to specify either the threshold value or the period of time during which symbol errors are counted (the window).

  The incorrect syntax is as follows:

  no ethernet oam lfm link-monitor symbol-period { threshold { high { *highFrames* | none } | low *lowFrames* } | window *period* }

  The correct syntax of the **no** version of this command is as follows:

  no ethernet oam lfm link-monitor symbol-period [threshold { high { *highFrames* | none } | low *lowFrames* } | window *period*]

  Using this syntax, you are not required to specify a threshold value or a time period. Both the **threshold** and **window** keywords are optional when you use the **no** form of this command.

  If you enter the **no ethernet oam lfm link-monitor symbol-period** command without specifying the threshold or window parameters, both the configured threshold and window attributes are removed. If you want to delete only one of the two settings, such as the threshold value or the period of time for monitoring symbol errors, you can use the keyword specific to the setting that you want to remove with the **no ethernet oam lfm link-monitor symbol-period** command.

- The description of the **ip route** command in the *JunosE Command Reference Guide A to M* incorrectly states that this command is available in VRF Configuration mode. The **ip route** command is available only in Global Configuration mode.
- In the *JunosE Command Reference Guide A to M*, the description of the **mpls ldp graceful-restart reconnect-time** command incorrectly states that the **no** version restores the default value of 120 seconds. The correct default value is 140 seconds.
- The *JunosE Command Reference Guide N to Z* omits the description of the **show memory-management protection** command, beginning with JunosE Release 7.1.0 in which it was introduced.

The **show memory-management protection** command is available in Privileged Exec mode. You can use this command to display the information about memory management protection of the router.

The syntax for this command is:

show memory-management protection [detail] [*filter*]

The command can be used only in the support mode and is not user configurable.

- In the *JunosE System Event Logging Reference Guide*, for the **radiusClient** event category, the Error field incorrectly includes the following errors:
  - Internal allocation error of base RADIUS server table
  - Invalid virtual router for user's context

  These errors were never implemented in any version of JunosE Software.

Appendix A

System Maximums

This appendix presents current system maximums for various E Series hardware configurations. An E Series router does not simultaneously support all maximum configurations.

For some entries, early field trial (EFT) values are presented in addition to supported values. These values have not been fully qualified by Juniper Networks and are mentioned only for field test purposes in this release. EFT values are enclosed within parentheses with an EFT designation; for example, (96,000 EFT).

Modules referred to in the tables are identified by their physical label. For module specifications, including their identifying labels, see *ERX Module Guide, Table 1, Module Combinations* and *E120 and E320 Module Guide, Table 1, Modules and IOAs*.

System Maximums for ERX310, ERX7xx, and ERX14xx	Section
General router values	*General System Maximums* on page 68
Physical layer values	*Physical and Logical Density Maximums* on page 69
Link layer values	*Link Layer Maximums* on page 72
Routing protocol and performance values	*Routing Protocol Maximums* on page 77
Policy and QoS values	*Policy and QoS Maximums* on page 80
Tunneling values	*Tunneling Maximums* on page 82
Subscriber management values	*Subscriber Management Maximums* on page 84

System Maximums for E120 and E320 Routers	Section
General router values	*General System Maximums* on page 86
Physical layer values	*Physical and Logical Density Maximums* on page 87
Link layer values	*Link Layer Maximums* on page 89
Routing protocol and performance values	*Routing Protocol Maximums* on page 95
Policy and QoS values	*Policy and QoS Maximums* on page 98
Tunneling values	*Tunneling Maximums* on page 102
Subscriber management values	*Subscriber Management Maximums* on page 104

ERX310, ERX7xx, and ERX14xx System Maximums

The following tables provide system maximums for the ERX310, ERX7xx, and ERX14xx routers.

General System Maximums

Table 1 lists some general system maximums for the ERX routers.

Table 1: General System Maximums

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Fabric size	10 Gbps	5 or 10 Gbps	10 Gbps	40 Gbps
Chassis per 7-foot rack	14	6	3	3
NTP clients	1000	1000	1000	1000
NTP servers	300	300	300	300
Sessions per chassis (simultaneous Telnet + FTP + SSH, in any combination)	30	30	30	30
Virtual routers per chassis	1000	1000	1000	1000
Virtual routers per line module ASIC	1000	1000	1000	1000
ICR Partitions per chassis	640	640	640	640
ICR Partitions per line module	64	64	64	64

Physical and Logical Density Maximums

Table 2 lists physical and logical density maximums for the ERX routers. The following notes are referred to in Table 2:

1. Wire rate indicates the port density that supports maximum (wire-rate) performance. Oversubscribed indicates the port density possible when you are willing to accept less than wire-rate performance by oversubscribing the available fabric bandwidth. The ERX310 and ERX1440 routers do not support oversubscription; port densities for these models indicate wire-rate performance.
2. When you pair the GE-2 or GE-HDE line module with the GE-2 SFP I/O module on the ERX1440 router, you can terminate up to 24 Gigabit Ethernet interfaces. Slots 2 and 4 on the ERX1440 router support two Gigabit Ethernet interfaces at wire rate; the remaining 10 slots support one Gigabit Ethernet interface at wire rate. On the ERX310 router, all four ports (active and redundant) are at wire rate.

   For more information about bandwidth and line-rate considerations for the GE-2 line module or the GE-HDE line module and their corresponding I/O modules on E Series routers, see *JunosE Physical Layer Configuration Guide, Chapter 5, Configuring Ethernet Interfaces*.
3. When you pair the GE-HDE line module with the GE-8 I/O module on the ERX1440 router, you can terminate up to 96 Gigabit Ethernet interfaces. Slots 2 and 4 on the ERX1440 router support two Gigabit Ethernet interfaces at wire rate; the remaining 10 slots support one Gigabit Ethernet interface at wire rate. On the ERX310 router, only two Gigabit Ethernet interfaces per slot are at wire rate; therefore, only four Gigabit Ethernet interfaces are at wire rate for the entire router.

   For more information about bandwidth and line-rate considerations for the GE-HDE line module and the GE-8 I/O module on E Series routers, see *JunosE Physical Layer Configuration Guide, Chapter 5, Configuring Ethernet Interfaces*.
4. The OC3/STM-1 GE/FE line module and OC3-2 GE APS I/O module combination does not support line rate for Gigabit Ethernet interfaces.

Table 2: Physical and Logical Density Maximums

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Physical density wire rate/oversubscribed (See Note 1 on page 69.)				
Channelized OC3 ports per chassis (cOC3 STM1 FO I/O modules)	8	16/20	32/48	48
Channelized OC12 ports per chassis (cOC12 STM4 FO I/O modules)	2	4/5	4/12	12
Channelized T3 ports per chassis (CT3/T3 12 I/O modules)	24	48/60	96/144	144
E3 (unchannelized) ports per chassis (CT3/T3 12 I/O modules)	24	48/60	96/144	144
Fast Ethernet (10/100) ports per chassis (FE-8 I/O and FE-8 SFP I/O modules)	16	32/40	32/96	96
Gigabit Ethernet ports per chassis (GE I/O modules)	2	4/5	4/12	12

Table 2: Physical and Logical Density Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Gigabit Ethernet ports per chassis (GE-2 SFP I/O modules) (See Note 2 on page 69.)	4	–	–	14/24
Gigabit Ethernet ports per chassis (GE-8 I/O modules) (See Note 3 on page 69.)	4/16	–	–	14/96
Gigabit Ethernet ports per chassis (OC3-2 GE APS I/O module) (See Note 4 on page 69.)	2	4/5	4/12	12
OC3/STM-1 ATM ports per chassis (OC3-4 I/O modules)	8	16/20	32/48	48
OC3/STM-1 ATM ports per chassis (OC3-2 GE APS I/O module)	4	10	24	24
OC3/STM-1 POS ports per chassis (OC3-4 I/O modules)	8	16/20	16/48	48
OC12/STM-4 ATM ports per chassis (OC12 STM4 I/O modules)	2	4/5	8/12	12
OC12/STM-4 POS ports per chassis (OC12 STM4 I/O modules)	2	4/5	4/12	12
OC48/STM16 POS ports per chassis (OC48 FRAME I/O modules); ERX1440 router only	–	–	–	2
T3 (unchannelized) ports per chassis (4xDS3 ATM I/O modules)	8	16/20	32/48	48
T3 (unchannelized) ports per chassis (CT3/T3 12 I/O modules)	24	48/60	96/144	144
Logical density per chassis				
Logical E1s per chassis	504	1260	3024	3024
Logical E3s per chassis	24	60	144	144
Logical fractional E1s (DS0) per chassis	4000	10,000	24,000	24,000
Logical fractional T1s (DS0) per chassis	4000	10,000	24,000	24,000
Logical OC3/STM1 per chassis	8	20	48	48
Logical OC12/STM4 per chassis	2	5	12	12
Logical OC48/STM16 per chassis (ERX1440 router only)	–	–	–	2
Logical T1s per chassis	672	1680	4032	4032
Logical T3s per chassis	24	60	144	144
Logical density per module combination (specified line module and all supported I/O modules)				
Logical E1s per cOCx/STMx F0 line module	252 63 per OC3/STM1	252 63 per OC3/STM1	252 63 per OC3/STM1	252 63 per OC3/STM1
Logical E3s per COCX-F3 line module	12	12	12	12

Table 2: Physical and Logical Density Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Logical fractional E1s (DS0) per cOCx/STMx F0 line module	2000 500 per OC3/STM1	2000 500 per OC3/STM1	2000 500 per OC3/STM1	2000 500 per OC3/STM1
Logical fractional T1s (DS0) per cOCx/STMx F0 line module	2000 500 per OC3/STM1	2000 500 per OC3/STM1	2000 500 per OC3/STM1	2000 500 per OC3/STM1
Logical fractional T1s (DS0) per CT3/T3-F0 line module	1992 166 per T3	1992 166 per T3	1992 166 per T3	1992 166 per T3
Logical fractional T3s (DS3) per COCX-F3 line module	12	12	12	12
Logical T1s per cOCx/STMx F0 line module	336 84 per OC3/STM1	336 84 per OC3/STM1	336 84 per OC3/STM1	336 84 per OC3/STM1
Logical T1s per CT3/T3-F0 line module	336 28 per T3	336 28 per T3	336 28 per T3	336 28 per T3
Logical T3s per COCX-F3 line module	12	12	12	12
Logical T3s per cOCx/STMx F0 line module	12 3 per OC3/STM1	12 3 per OC3/STM1	12 3 per OC3/STM1	12 3 per OC3/STM1
Logical T3s per CT3/T3-F0	12	12	12	12
Logical T3s per OCx/STMx/DS3-ATM line module with 4xDS3 ATM I/O module	4	4	4	4

Link Layer Maximums

Table 3 lists link layer maximums for the ERX routers. The following notes are referred to in Table 3:

1. The ERX1440 router supports a maximum of 48,000 interface columns of all types combined. You can use either all dynamic interfaces or a combination of dynamic and static interfaces to achieve this maximum. For bridged Ethernet, IP network, and PPP interfaces, the ERX1440 router supports a maximum of 32,000 static major interfaces. Although the ERX1440 router supports a maximum of 48,000 static major interfaces for PPPoE, the PPPoE static limit is enforced at the subinterface level, which has a limit of 32,000.

   The ERX705, ERX710, and ERX1410 routers support a maximum of 32,000 interfaces of all types combined; the ERX310 router supports a maximum of 16,000 interfaces of all types combined. For these routers, the interfaces can be any combination of dynamic or static.

   The JunosE Software supports up to 10,000 PPP interfaces with EAP authentication negotiation configured. Performance and scalability is unchanged when EAP is not configured.

2. The total maximum number of Ethernet subinterfaces that can be active at any one time on an ERX310 router, an ERX7xx router, or an ERX14xx router is limited by the number of slots per chassis. Of this total, you can configure all single-tagged VLAN subinterfaces, all double-tagged S-VLAN subinterfaces, or a combination of both VLAN subinterfaces and S-VLAN subinterfaces to achieve this maximum.

Table 3: Link Layer Maximums

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
ARP entries per line module				
Dynamic ARP entries	32,768	32,768	32,768	32,768
Static ARP entries	32,768	32,768	32,768	32,768
Total ARP entries	32,768	32,768	32,768	32,768
ATM bulk configuration VC ranges per chassis	300	300	300	300
ATM bulk configuration VC ranges per line module	300	300	300	300
ATM bulk configuration total VCs per chassis	64,000	160,000	384,000	384,000
ATM bulk configuration total VCs per line module				
OCx/STMx/DS3-ATM	32,000	32,000	32,000	32,000
OC3/STM1 GE/FE	32,000	32,000	32,000	32,000
ATM bulk configuration overriding profile assignments per chassis	100	100	100	100
ATM VCs per chassis (active/configured)	16,000/32,000	32,000/64,000	32,000/64,000	48,000/96,000

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
ATM VCs per line module				
OCx/STMx/DS3-ATM (active/configured)	8000/16,000	8000/16,000	8000/16,000	8000/16,000
OC3/STM1 GE/FE (active/configured)	8000/16,000	8000/16,000	8000/16,000	8000/16,000
ATM VCs per port				
OCx/STMx/DS3-ATM (active/configured)	8000/16,000	8000/16,000	8000/16,000	8000/16,000
OC3/STM1 GE/FE (active/configured)	8000/16,000	8000/16,000	8000/16,000	8000/16,000
ATM VC classes per chassis	100	100	100	100
ATM VP/VC addresses per line module				
OCx/STMx/DS3-ATM	20-bit	20-bit	20-bit	20-bit
OC3/STM1 GE/FE	20-bit	20-bit	20-bit	20-bit
ATM VP tunnels per port, all supported modules	256	256	256	256
Bridged Ethernet interfaces per chassis (See Note 1 on page 72.)	16,000	32,000	32,000	48,000
Bridged Ethernet interfaces per line module				
OCx/STMx/DS3-ATM	8192	8192	8192	8192
OC3/STM-1 GE/FE	8192	8192	8192	8192
Dynamic interfaces				
Active autosensed dynamic interface columns per chassis over static or dynamic (bulk) ATM1483 subinterfaces	16,000	32,000	32,000	48,000
Ethernet 802.3ad Link Aggregation				
Links per LAG (bundle)	8	8	8	8
LAGs (bundles) per chassis	64	64	64	64
Ethernet S-VLANs per chassis (See Note 2 on page 72.)	32,768	81,920	96,000	96,000
Ethernet S-VLANs per I/O module				
FE-8 I/O and FE-8 SFP I/O	16,384	16,384	16,384	16,384
GE I/O	16,384	16,384	16,384	16,384
GE-2 SFP I/O	16,384	–	–	16,384
GE-8 I/O	16,384	–	–	16,384
OC3-2 GE APS I/O	16,384	16,384	16,384	16,384

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Ethernet VLANs per chassis (See Note 2 on page 72.)	32,768	81,920	96,000	96,000
Ethernet VLANs per I/O module (no more than 4096 VLANs per port)				
FE-8 I/O and FE-8 SFP I/O	8192	8192	8192	8192
GE I/O	4096	4096	4096	4096
GE-2 SFP I/O	8192	–	–	8192
GE-8 I/O	16,384	–	–	16,384
OC3-2 GE APS I/O	4096	4096	4096	4096
Ethernet VLAN bulk configuration VLAN ranges per chassis	300	300	300	300
Ethernet VLAN bulk configuration VLAN ranges per line module	300	300	300	300
Ethernet VLAN overriding profile assignments per chassis	200	200	200	200
Ethernet VRRP VRIDs per line module ASIC	800	800	800	800
Frame Relay virtual circuits per chassis	2000	5000	12,000	12,000
Frame Relay virtual circuits per line module				
COCX-F3	1000	1000	1000	1000
cOCx/STMx F0	1000	1000	1000	1000
OC48 (ERX1440 router only)	–	–	–	1000
Frame Relay virtual circuits per port				
COCX-F3	1000	1000	1000	1000
cOCx/STMx F0	1000	1000	1000	1000
OC48 (ERX1440 router only)	–	–	–	1000
HDLC interfaces per chassis	4000	10,000	24,000	24,000
HDLC interfaces per line module				
COCX-F3	12	12	12	12
cOCx/STMx F0	2000	2000	2000	2000
CT3/T3 F0	1992	1992	1992	1992
OCx/STMx/DS-3 ATM	8000	8000	8000	8000
OCx/STMx POS	4	4	4	4
OC48 (ERX1440 router only)	–	–	–	1

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
MLFR bundles per chassis	5000	5000	5000	5000
MLFR bundles per line module	Bundles per line module are limited only by the availability of interface columns on the module. Because a bundle requires at least one interface column, the number of bundles cannot exceed the number of interface columns.			
MLPPP bundles per chassis	12,000	12,000	12,000	12,000
MLPPP bundles per line module	The maximum number of MLPPP bundles supported per line module is the lesser of the maximum number of MLPPP bundles supported per chassis or of the maximum number of interfaces supported on the line module. For more information, see the *JunosE Link Layer Configuration Guide*.			
PPP interfaces per chassis (See Note 1 on page 72.)	16,000	32,000	32,000	48,000
PPP interfaces per line module				
COCX-F3	12	12	12	12
cOCx/STMx FO	2000	2000	2000	2000
GE/FE	8000	8000	8000	8000
GE-2	8000	–	–	8000
GE-HDE	8000	–	–	8000
OCx/STMx/DS-3 ATM	8000	8000	8000	8000
OC3/STM-1 GE/FE	8000	8000	8000	8000
OCx/STMx POS	4	4	4	4
OC48 (ERX1440 router only)	–	–	–	1
PPP packet logging				
Aggregate dynamic and static PPP interfaces for which you can log PPP packets per chassis	32	32	32	32
PPPoE service name tables				
PPPoE service name tables per chassis	16	16	16	16
Service name tags per PPPoE service name table (including one empty service name tag)	17	17	17	17
PPPoE subinterfaces				
Subinterfaces per chassis (See Note 1 on page 72.)	16,000	32,000	32,000	48,000
Subinterfaces per GE/FE line module	8000	8000	8000	8000
Subinterfaces per GE-2 line module	8000	–	–	8000

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Subinterfaces per GE-HDE line module	8000	–	–	8000
Subinterfaces per OCx/STMx/DS-3 ATM line module	8000	8000	8000	8000
Subinterfaces per OC3/STM-1 GE/FE line module	8000	8000	8000	8000
Transparent bridging and VPLS				
Bridge groups or VPLS instances per chassis	1024	1024	1024	1024
Bridge interfaces per line module in bridge groups or VPLS instances	8000	8000	8000	8000
Bridge interfaces per chassis in bridge groups or VPLS instances	16,000	32,000	32,000	32,000
Learned MAC address entries combined for all bridge groups and VPLS instances on a chassis	64,000	64,000	64,000	64,000

Routing Protocol Maximums

Table 4 lists routing protocol maximums for the ERX routers. The following notes are referred to in Table 4:

1. 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 three 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.
2. The ERX1440 router supports a maximum of 48,000 interfaces of all types combined. You can use either all dynamic interfaces or a combination of dynamic and static interfaces to achieve this maximum. The ERX1440 router supports a maximum of 32,000 static PPP/PPPoE interfaces and a maximum of 36,500 static IP network interfaces. Bridged Ethernet does not enforce a limit so IP interfaces created on Bridged Ethernet can scale to the IP maximum of 36,500.

   The ERX705, ERX710, and ERX1410 routers support a maximum of 32,000 IP network interfaces; the ERX310 router supports a maximum of 16,000 IP network interfaces. For all these models, the interfaces can be any combination of dynamic or static.
3. These values are subject to limitations on available SRP module memory, which varies according to your router configuration.
4. Depending on your configuration, the router may support more routing table entries or fewer routing table entries than this value. In any case, you can choose to limit the number of routes that can be added to the routing table on a per-VR or per-VRF basis by means of the **maximum routes** command.
5. The maximum number of ANCP adjacencies can be scaled over a maximum of 100 virtual routers. Fewer ANCP adjacencies can be scaled in configurations with more than 100 virtual routers.
6. This maximum is not valid for Frame Relay. The Frame Relay maximum is 1000 circuits over MPLS per line module, because only 1000 Frame Relay DLCIs are permitted per line module.
7. On the ERX1440 router, you can achieve 32,767 total Martini circuits over ATM or Ethernet interfaces. For all routers, the total Martini can be any combination of external inter-router circuits and internal circuits (local cross-connects).
8. There is no per-VR limit; all multicast routes can be on a single VR or present across multiple VRs.
9. The maximum number of interfaces can be achieved by any combination; for example, two streams each being replicated to 32,768 interfaces; 16,384 streams each being replicated four times; or any other combination.

10. Dynamic values represent typical limits that vary depending on configuration details and actual dynamic behavior. For dynamic values only, multiple server modules (SMs) in a chassis can improve the values as long as the multiple server modules are online and the number of virtual routers configured with NAT is greater than or equal to the number of server modules. If a server module fails, the load is redistributed to the remaining server modules, with a consequent reduction in aggregate capacity.
11. Static and dynamic translations occupy the same table; therefore, the number of static translation entries present in the table reduces the room for dynamic entries.

Table 4: Routing Protocol Maximums

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
BFD				
Sessions per line module	50	50	50	50
ECMP maximum paths to a destination				
BGP, IS-IS, MPLS, OSPF, RIP	16	16	16	16
IPv4 forwarding table entries (See Note 1 on page 77.)				
Chassis with only ASIC modules	1,048,576	1,048,576	1,048,576	1,048,576
IP network interfaces (IPv4 and IPv6)				
Per chassis (See Note 2 on page 77.)	16,000	32,000	32,000	48,000
Per line module ASIC	8000	8000	8000	8000
IPv4 routing protocol scaling and peering densities (See Note 3 on page 77.)				
Routing table entries (See Note 4 on page 77.)	500,000	500,000	500,000	500,000
ANCP Adjacency Scaling (See Note 5 on page 77.)	5000	5000	5000	5000
BGP-4 peering sessions	1000	1000	1000	1000
BGP-4 routes (NLRI)	1,500,000	1,500,000	1,500,000	1,500,000
IP next hops (egress FECs) on router with ASIC modules (used to represent the IP addresses of next-hop routers on Ethernet interfaces)	1,000,000	1,000,000	1,000,000	1,000,000
MPLS next hops (egress FECs) on router with ASIC modules only	500,000	500,000	500,000	500,000
MPLS forwarding entries	64,000	64,000	64,000	64,000
IS-IS adjacencies	150	150	150	150
IS-IS routes	20,000	20,000	20,000	20,000
MPLS LDP LSPs	10,000	10,000	10,000	10,000
MPLS RSVP-TE LSPs	10,000	10,000	10,000	10,000

Table 4: Routing Protocol Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
OSPF adjacencies	1000	1000	1000	1000
OSPF routes	25,000	25,000	25,000	25,000
IPv6 routing table entries (See Note 3 on page 77.)	50,000	50,000	50,000	50,000
J-Flow statistics				
J-Flow–enabled VRs and VRFs, in any combination	16	16	16	16
Sampled interfaces per VR or VRF	32	32	32	32
Total sampled Interfaces per chassis	512	512	512	512
Martini circuits for layer 2 services over MPLS				
Total Martini circuits per line module (See Note 6 on page 77.)	8000	8000	8000	8000
Total Martini circuits per chassis (See Note 7 on page 77.)	16,000	16,000	16,000	32,767
External Martini circuits per chassis	16,000	16,000	16,000	32,767
Internal Martini circuits (local cross-connects) per chassis	16,000	16,000	16,000	32,767
Mobile IP bindings per chassis	–	–	–	48,000
Multicast routes (IPv4 and IPv6)				
Forwarding entries [(S,G) pairs] per chassis (See Note 8 on page 77.)	16,384	16,384	16,384	16,384
Outgoing interfaces per chassis (See Note 9 on page 77.)	65,536	65,536	65,536	65,536
Network Address Translation (NAT)				
Static translations (simple or extended) per chassis	96,000	96,000	96,000	96,000
Dynamic simple translations (NAT) per SM (See Notes 10 and 11 on page 78.)	400,000	400,000	400,000	400,000
Dynamic extended translations (NAPT) per SM (See Notes 10 and 11 on page 78.)	200,000	200,000	200,000	200,000
Response Time Reporter simultaneous operations per VR	500	500	500	500
VRRP VRIDs per line module ASIC	See *Ethernet VRRP VRIDs per line module ASIC* on page 74.			

Policy and QoS Maximums

Table 5 lists policy and QoS maximums for the ERX routers. The following notes are referred to in Table 5:

1. The OC48 line module supports only 131,071 entries. The GE-2 and GE-HDE line modules support only 65,535 entries.
2. For line modules other than the GE-2, GE-HDE, and OC48/STM16 line modules, the router supports two sizes of policies: 8127 policies, each with a maximum of 32 classifiers, and 16,255 policies, each with a maximum of 16 classifiers. A combination of the two sizes of policies is also supported, in which case the total number of policies is between 8127 and 16,255, depending on the actual configuration.
3. The GE-2, GE-HDE, and OC48/STM16 line modules support CAM classifiers instead of hardware policy assignments. For most configurations, each classifier entry in a policy consumes one CAM entry. However, a policy that has only the default classifier consumes no CAM resources. Policies that use CAM hardware classifiers consume one interface attachment resource, regardless of the number of classifier entries in a policy.

Table 5: Policy and QoS Maximums

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
QoS queues per ASIC line module	49,000	49,000	49,000	49,000
QoS profiles configurable per chassis	1000	1000	1000	1000
QoS profile attachments per chassis	96,000	96,000	96,000	96,000
QoS profile attachments per ASIC line module	16,000	16,000	16,000	16,000
QoS shapers per line module	64,000	64,000	64,000	64,000
Classification rules per policy	512	512	512	512
Policy classification (CLACL) entries per line module (See Note 1 on page 80.)	256,000	256,000	256,000	256,000
Unique hardware policy assignments per line module for modules other than the GE-2, GE-HDE, and OC48/STM16 (See Note 2 on page 80.)	8127/16,255	8127/16,255	8127/16,255	8127/16,255

Table 5: Policy and QoS Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
CAM entries (See Note 3 on page 80.)				
GE-2	64,000	–	–	64,000
GE-HDE	64,000	–	–	64,000
OC48/STM16	–	–	–	128,000
Policy egress interface attachments per ASIC line module				
Combined IP and IPv6 interface attachments	8191	8191	8191	8191
Combined ATM, Frame Relay, GRE, L2TP (LNS only), MPLS, and VLAN interface attachments	8191	8191	8191	8191
Policy ingress interface attachments per ASIC line module				
Combined IP and IPv6 interface attachments on GE-2, GE-HDE, and OC-48/STM16 line modules	16,383	–	–	16,383
Combined IP and IPv6 interface attachments on all other line modules	16,000	16,000	16,000	16,000
Combined ATM, Frame Relay, GRE, L2TP (LNS only), MPLS, and VLAN interface attachments	8191	8191	8191	8191
Rate limiters				
Egress per ASIC line module	24,575	24,575	24,575	24,575
Ingress per ASIC line module	24,575	24,575	24,575	24,575
Policy statistics blocks				
Egress per ASIC line module	256,000	256,000	256,000	256,000
Ingress per ASIC line module	256,000	256,000	256,000	256,000
Software lookup blocks				
Per ASIC line module	16,383	16,383	16,383	16,383
Secure policies (for packet mirroring)				
Per ASIC line module	1022	1022	1022	1022
Per chassis	2400	2400	2400	2400

Tunneling Maximums

Table 6 lists tunneling maximums for the ERX routers. The following notes are referred to in Table 6:

1. The SM supports any combination of DVMRP, GRE, and L2TP tunnels up to a maximum of 8000 tunnels; however, no more than 4000 tunnels can be DVMRP or GRE tunnels in any combination. The ISM supports any combination of DVMRP, GRE, and L2TP tunnels over IPSec, up to a maximum of 5000 tunnels; however, no more than 4000 tunnels can be DVMRP or GRE tunnels.
2. You can have no more than 8000 L2TP/IPSec sessions per chassis.
3. For more information about supported L2TP sessions and tunnels, see *JunosE Broadband Access Configuration Guide, Chapter 11, L2TP Overview*.

Table 6: Tunneling Maximums

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
DVMRP (IP-in-IP) tunnels per chassis	4000	4000	4000	4000
DVMRP (IP-in-IP) tunnels per line module (See Note 1 on page 82.)				
GE-2 with shared tunnel-server ports provisioned	4000	–	–	4000
GE-HDE with shared tunnel-server ports provisioned	4000	–	–	4000
IPSec Service Module (DVMRP/IPSec tunnels)	4000	4000	4000	4000
Service Module (SM)	4000	4000	4000	4000
GRE tunnels per chassis	4000	4000	4000	4000
GRE tunnels per line module (See Note 1 on page 82.)				
GE-2 with shared tunnel-server ports provisioned	4000	–	–	4000
GE-HDE with shared tunnel-server ports provisioned	4000	–	–	4000
IPSec Service Module (GRE/IPSec tunnels)	4000	4000	4000	4000
Service Module (SM)	4000	4000	4000	4000
IPSec manual secure tunnels per chassis	256	256	256	256
IPSec transform sets per chassis	1000	1000	1000	1000
IPSec transforms per transform set	6	6	6	6
IPSec tunnels per chassis	10,000	10,000	10,000	20,000
IPSec tunnels per IPSec Service Module	5000	5000	5000	5000

Table 6: Tunneling Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
L2TP sessions per chassis (See Notes 2 and 3 on page 82.)	16,000	16,000	16,000	32,000
L2TP sessions per line module (See Notes 1 and 3 on page 82.)				
GE-2 with shared tunnel-server ports provisioned	8000	–	–	8000
GE-HDE with shared tunnel-server ports provisioned	8000	–	–	8000
IPSec Service Module (ISM; L2TP/IPSec sessions)	5000	5000	5000	5000
Service Module (SM)	16,000	16,000	16,000	16,000
L2TP tunnels per chassis	8000	8000	8000	8000
L2TP tunnels per line module (See Notes 1 and 3 on page 82.)				
GE-2 with shared tunnel-server ports provisioned	8000	–	–	8000
GE-HDE with shared tunnel-server ports provisioned	8000	–	–	8000
IPSec Service Module (L2TP/IPSec tunnels)	5000	5000	5000	5000
Service Module	8000	8000	8000	8000

Subscriber Management Maximums

Table 7 lists subscriber management maximums for the ERX routers. The following notes are referred to in Table 7:

1. DHCP relay proxy maintains a list of active DHCP clients up to a maximum of 100,000 clients per chassis for all virtual routers. DHCP relay does not maintain a list of DHCP clients.

   DHCP relay proxy is notified of DHCP client deletions and subsequently deletes the client's host routes. In contrast, DHCP relay is not notified of DHCP client deletions, so the host routes for deleted clients remain in DHCP relay until you permanently delete them with the **set dhcp relay discard-access-routes** command. A maximum of 100,000 host routes for DHCP clients can be stored for all DHCP relay and DHCP relay proxy instances (that is, for all virtual routers).

2. The ERX1440 router supports a maximum of 48,000 interface columns of all types combined. You can use either all dynamic interfaces or a combination of dynamic and static interfaces to achieve this maximum. For bridged Ethernet, IP network, and PPP interfaces, the ERX1440 router supports a maximum of 32,000 static major interfaces. Although the ERX1440 router supports a maximum of 48,000 static major interfaces for PPPoE, the PPPoE static limit is enforced at the subinterface level, which has a limit of 32,000.

   The ERX705, ERX710, and ERX1410 routers support a maximum of 32,000 interfaces of all types combined; the ERX310 router supports a maximum of 16,000 interfaces of all types combined. For these routers, the interfaces can be any combination of dynamic or static.

   The JunosE Software supports up to 10,000 PPP interfaces with EAP authentication negotiation configured. Performance and scalability is unchanged when EAP is not configured.

3. For DHCPv6 local server, up to 32,000 subscribers and clients are supported on PPP/ATM and PPPoE/ATM with dynamic interfaces. Interface flapping tests have been qualified for 8000 subscribers and interfaces.

Table 7: Subscriber Management Maximums

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
DHCP external server clients (per chassis for all virtual routers; and per virtual router) (See Note 1 on page 84.)	100,000	100,000	100,000	100,000
DHCP local server (See Note 2 on page 84.)				
Client bindings per chassis	96,000	96,000	96,000	96,000
Client interfaces per chassis	16,000	32,000	32,000	48,000
Local address pools per virtual router	4000	4000	4000	4000
IP addresses per local address pool	32,000	32,000	32,000	32,000

Table 7: Subscriber Management Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
DHCPv6 local server				
Clients (See Note 3 on page 84.)	32,000	32,000	32,000	32,000
DHCP relay and relay proxy client (See Notes 1 and 2 on page 84.)				
DHCP client host routes for DHCP relay and DHCP relay proxy combined (per chassis for all virtual routers; and per virtual router)	100,000	100,000	100,000	100,000
DHCP relay proxy clients (per chassis for all virtual routers; and per virtual router)	100,000	100,000	100,000	100,000
Interfaces (per chassis for all virtual routers; and per virtual router)	16,000	32,000	32,000	48,000
Local authentication server				
Local user databases per chassis	100	100	100	100
Users per local user database	100	100	100	100
Users for all local user databases	100	100	100	100
RADIUS requests				
Concurrent RADIUS authentication requests	4000	4000	4000	32,000
Concurrent RADIUS accounting requests	4000	4000	4000	96,000
RADIUS route-download server downloaded routes per chassis	32,000	32,000	32,000	32,000
Service Manager				
Service definitions	2048	2048	2048	2048
Service sessions (active)	131,072	131,072	131,072	131,072
Active subscriber sessions	16,000	32,000	32,000	48,000
SRC Software and SDX Software				
COPS client instances	200	200	200	200
SRC clients	200	200	200	200
SRC interfaces	16,000	32,000	32,000	48,000
Subscriber interfaces (See Note 2 on page 84.)				
Dynamic subscriber interfaces per chassis	16,000	32,000	32,000	48,000
Dynamic subscriber interfaces per line module	8000	8000	8000	8000
Static subscriber interfaces per chassis	16,000	32,000	32,000	48,000
Static subscriber interfaces per line module	8000	8000	8000	8000

E120 and E320 System Maximums

The following tables provide system maximums for the E120 router and the E320 router.

General System Maximums

Table 8 lists some general system maximums for the E120 router and the E320 router. The following notes are referred to in Table 8:

1. The maximum number applies to any combination of VRs and VRFs. The number of VRs and VRFs that you can configure depends on your configuration. You cannot achieve the maximum number if each VR and VRF instance is running a routing protocol.
2. The maximum of 3000 VRs and VRFs can be achieved only with the SRP-120 and SRP-320 modules, which have 4 GB of memory. The limits cannot be achieved with the SRP-100 module, which has 2 GB of memory.

Table 8: General System Maximums

Feature	E120	E320
Fabric size	120 Gbps	100 Gbps/320 Gbps
Chassis per 7-foot rack	6	3
NTP clients	1000	1000
NTP servers	300	300
Sessions per chassis (simultaneous Telnet + FTP + SSH, in any combination)	30	30
Virtual routers and VRFs per chassis, combined (See Notes 1 and 2 on page 86.)	3000	3000
Virtual routers and VRFs per line module, combined (See Notes 1 and 2 on page 86.)	3000	3000
ICR Partitions per chassis	640	640
ICR Partitions per line module	64	64

Physical and Logical Density Maximums

Table 9 lists physical and logical density maximums for the E120 router and the E320 router. The following notes are referred to in Table 9:

1. Wire rate indicates the port density that supports maximum (wire-rate) performance. Oversubscribed indicates the port density possible if you are willing to accept less than wire-rate performance by oversubscribing the available fabric bandwidth.
2. With a 120 Gbps configuration on the E120 router, you can install up to 6 combinations of ES2 10G Uplink LMs, ES2 10G LMs, or ES2 10G ADV LMs in slots numbered 0-5. You can install a maximum of 6 active ports and 6 redundant ports at any time.

   With a 100 Gbps fabric configuration on the E320 router, you must install the ES2 10G Uplink LM or the ES2 10G LM in either of the E320 router turbo slots (2 and 4). When the ES2 10G Uplink LM or the ES2 10G LM is installed in slot 2 or slot 4, you cannot install another line module in slot 3 or slot 5. In this case, you can only install the ES2 4G LM in slots 0–1 and 6–11; therefore, the maximum number of ports and the forwarding performance per chassis is reduced for the IOAs that pair with the ES2 4G LM.

   With a 320 Gbps fabric configuration on the E320 router, you can install up to 12 combinations of ES2 10G Uplink LMs, ES2 10G LMs, or ES2 10G ADV LMs in slots numbered 0-5 and 11-16. You can install a maximum of 12 active ports and 12 redundant ports at any time.

Table 9: Physical and Logical Density Maximums

Feature	E120	E320
Physical density wire rate/oversubscribed (See Note 1 on page 87.)		
10-Gigabit Ethernet ports per chassis (ES2-S1 10GE IOA)	6	12
10-Gigabit Ethernet ports per chassis (ES2-S2 10GE PR IOA) (See Note 2 on page 87.)	6 + 6	12 + 12
Gigabit Ethernet ports per chassis (ES2-S1 GE-4 IOAs)	24	48
Gigabit Ethernet ports per chassis (ES2-S1 GE-8 IOAs) (See Note 2 on page 87.)	96	192
Gigabit Ethernet ports per chassis (ES2-S3 GE-20 IOA) (See Note 2 on page 87.)	120	240
OC3/STM-1 ATM ports per chassis (ES2-S1 OC3-8 STM1 ATM IOAs)	96	192
OC12/STM-4 ATM ports per chassis (ES2-S1 OC12-2 STM4 ATM IOAs)	24	48
OC12/STM-4 POS ports per chassis (ES2-S1 OC12-2 STM4 POS IOAs)	24	48
OC48/STM16 ports per chassis (ES2-S1 OC48 STM16 POS IOAs)	6	12

Table 9: Physical and Logical Density Maximums Table continued

Feature	E120	E320
Logical density per chassis		
Logical OC3/STM1 per chassis	96	192
Logical OC12/STM4 per chassis	24	48
Logical OC48/STM16 per chassis	6	12

Link Layer Maximums

Table 10 lists link layer maximums for the E120 router and the E320 router. The following notes are referred to in Table 10:

1. On the ES2 10G LM, ES2 10G ADV LM, or ES2 10 G Uplink LM, you can have configurations with up to 100,000 static entries that support 100,000 DHCP relay proxy clients. You can have an additional 28,000 static or dynamic entries for network resources, such as RADIUS and DHCP servers. However, the total number of dynamic entries in the ARP table is still restricted to a maximum of 32,768 per line module.
2. On the E120 router, the SRP-120 and the SRP-320 support a maximum of 64,000 interfaces.

   On the E320 router, the SRP-320 supports a maximum of 96,000 interfaces. The SRP-100 supports a maximum of 64,000 interfaces.
3. The E120 router supports a maximum of 64,000 interface columns of all types combined. The E320 router supports a maximum of 96,000 interface columns of all types combined. You can use all dynamic interfaces, or all static interfaces, or a combination of dynamic and static interfaces to achieve this maximum.

   The JunosE Software supports up to 10,000 PPP interfaces with EAP authentication negotiation configured. Performance and scalability is unchanged when EAP is not configured.
4. The E120 router supports a maximum of 64,000 Ethernet subinterfaces that can be active at any one time. The E320 router supports a maximum of 96,000 Ethernet subinterfaces that can be active at any one time. Of this total, you can configure all single-tagged VLAN subinterfaces, all double-tagged S-VLAN subinterfaces, or a combination of both VLAN subinterfaces and S-VLAN subinterfaces to achieve this maximum.
5. The E120 router and the E320 router support 16,384 VLAN subinterfaces per slot on the ES2 4G LM and the ES2 10G LM, and 32,768 VLAN subinterfaces per slot on the ES2 10G ADV LM. On the E120 router, a maximum of 64,000 VLAN subinterfaces is supported per chassis. On the E320 router, a maximum of 96,000 VLAN subinterfaces is supported per chassis. You can use all dynamic interfaces, or all static interfaces, or a combination of dynamic and static interfaces to achieve this maximum.
6. For all LMs, no more than 16,384 S-VLANs are supported per port. The ES2 10G ADV LM supports 32,768 S-VLANs per module. All other LMs support only 16,384 S-VLANs per module.
7. For all LMs, no more than 4096 VLANs are supported per port. The ES2 10G ADV LM supports 32,768 VLANs per module. All other LMs support only 16,384 VLANs per module.
8. No more than 8192 VLAN major interfaces are supported per line module.

Table 10: Link Layer Maximums

Feature	E120	E320
ARP entries per line module		
Dynamic entries per LM	32,768	32,768
Static entries per ES2 4G LM	32,768	32,768

Table 10: Link Layer Maximums Table continued

Feature	E120	E320
Static entries per ES2 10G LM, ES2 10G ADV LM, or ES2 10G Uplink LM (See Note 1 on page 89.)	128,000	128,000
Total entries per ES2 4G LM	32,768	32,768
Total entries per ES2 10G LM, ES2 10G ADV LM, or ES2 10G Uplink LM (See Note 1 on page 89.)	128,000	128,000
ATM bulk configuration VC ranges per chassis	300	1025
ATM bulk configuration VC ranges per line module	300	1025
ATM bulk configuration total VCs per chassis	192,000	384,000
ATM bulk configuration total VCs per line module		
ES2 4G LM and OCx/STMx ATM IOA	32,000	32,000
ATM bulk configuration overriding profile assignments per chassis	100	100
ATM VCs per chassis (See Note 2 on page 89.)	64,000	96,000
ATM VCs per line module		
ES2 4G LM and OCx/STMx ATM IOA	16,000	16,000
ATM VCs per port		
ES2 4G LM and OCx/STMx ATM IOA	16,000	16,000
ATM VC classes per chassis	100	100
ATM VP/VC addresses per line module		
ES2 4G LM and OCx/STMx ATM IOA	24-bit	24-bit
ATM VP tunnels per port, all supported modules	256	256
Bridged Ethernet interfaces per chassis (See Notes 2 and 3 on page 89.)	64,000	96,000

Table 10: Link Layer Maximums Table continued

Feature	E120	E320
Bridged Ethernet interfaces per line module (OCx/STMx ATM)	16,000	16,000
Dynamic interfaces		
Active autosensed dynamic interface columns per chassis over static or dynamic (bulk) ATM1483 subinterfaces (See Note 2 on page 89.)	64,000	96,000
Ethernet 802.3ad Link Aggregation		
Links per LAG (bundle)	8	8
LAGs (bundles) per chassis	64	64
Ethernet S-VLANs per chassis (See Notes 2, 4, and 5 on page 89.)	64,000	96,000
Ethernet S-VLANs per IOA (See Note 6 on page 89.)		
ES2-S1 GE-4 IOA (with ES2 4G LM)	16,384	16,384
ES2-S1 GE-8 IOA (with ES2 4G LM or ES2 10G LM)	16,384	16,384
ES2-S1 GE-8 IOA (with ES2 10G ADV LM)	32,768	32,768
ES2-S1 10GE IOA (with ES2 4G LM)	16,384	16,384
ES2-S2 10GE PR IOA (with ES2 10G LM or ES2 10G Uplink LM)	16,384	16,384
ES2-S2 10GE PR IOA (with ES2 10G ADV LM)	32,768	32,768
ES2-S3 GE-20 IOA (with ES2 10G LM)	16,384	16,384
ES2-S3 GE-20 IOA (with ES2 10G ADV LM)	32,768	32,768
Ethernet VLANs per chassis (See Notes 2, 4, and 5 on page 89.)	64,000	96,000
Ethernet VLANs per IOA (See Note 7 on page 89.)		
ES2-S1 GE-4 IOA (with ES2 4G LM) (See Note 5 on page 89.)	16,384	16,384
ES2-S1 GE-8 IOA (with ES2 4G LM or ES2 10G LM) (See Note 5 on page 89.)	16,384	16,384

Table 10: Link Layer Maximums Table continued

Feature	E120	E320
ES2-S1 GE-8 IOA (with ES2 10G ADV LM) (See Note 5 on page 89.)	32,768	32,768
ES2-S1 10GE IOA (with ES2 4G LM) (See Note 5 on page 89.)	16,384	16,384
ES2-S2 10GE PR IOA (with ES2 10G LM, ES2 10G ADV LM, or ES2 10G Uplink LM) (See Note 5 on page 89.)	4096	4096
ES2-S3 GE-20 IOA (with ES2 10G LM)	16,384	16,384
ES2-S3 GE-20 IOA (with ES2 10G ADV LM)	32,768	32,768
Ethernet VLAN major interfaces over Bridged Ethernet Interfaces, per IOA (See Note 8 on page 89.)		
ES2-S1 GE-4 IOA (with ES2 4G LM)	8192	8192
ES2-S1 GE-8 IOA (with ES2 4G LM, ES2 10G LM, or ES2 10G ADV LM)	8192	8192
ES2-S1 10GE IOA (with ES2 4G LM)	8192	8192
ES2-S2 10GE PR IOA (with ES2 10G LM, ES2 10G ADV LM, or ES2 10G Uplink LM)	4096	4096
ES2-S3 GE-20 IOA (with ES2 10G LM or ES2 10G ADV LM)	8192	8192
Ethernet VLAN bulk configuration VLAN ranges per chassis	1000	1000
Ethernet VLAN bulk configuration VLAN ranges per line module	500	500
Ethernet VLAN overriding profile assignments per chassis	200	200
Ethernet VRRP VRIDs per line module	800	800
HDLC interfaces per chassis	24,000	24,000
HDLC interfaces per line module	8000	8000
MLPPP bundles per chassis	12,000	12,000

Table 10: Link Layer Maximums Table continued

Feature	E120	E320
MLPPP bundles per line module	The maximum number of MLPPP bundles supported per line module is the lesser of the maximum number of MLPPP bundles supported per chassis or of the maximum number of interfaces supported on the line module. For more information, see the *JunosE Link Layer Configuration Guide*.	
PPP major interfaces per chassis (See Notes 2 and 3 on page 89.)	64,000	96,000
PPP major interfaces per line module (ignoring physical interface constraints)		
ES2 4G LM	16,000	16,000
ES2 10G LM	16,000	16,000
ES2 10G ADV LM	32,000	32,000
PPP subinterfaces per chassis (See Notes 2 and 3 on page 89.)	64,000	96,000
PPP subinterfaces per line module (ignoring physical interface constraints)		
ES2 4G LM	16,000	16,000
ES2 10G LM	16,000	16,000
ES2 10G ADV LM	32,000	32,000
PPP packet logging		
Aggregate dynamic and static PPP interfaces for which you can log PPP packets per chassis	32	32
PPPoE service name tables		
PPPoE service name tables per chassis	16	16
Service name tags per PPPoE service name table (including one empty service name tag)	17	17
PPPoE subinterfaces per chassis (See Notes 2 and 3 on page 89.)	64,000	96,000
PPPoE subinterfaces per line module		
ES2 4G LM	16,000	16,000
ES2 10G LM	16,000	16,000
ES2 10G ADV LM	32,000	32,000

Table 10: Link Layer Maximums Table continued

Feature	E120	E320
Transparent bridging and VPLS		
Bridge groups or VPLS instances per chassis	1024	1024
Bridge interfaces per line module in bridge groups or VPLS instances	8000	8000
Bridge interfaces per chassis in bridge groups or VPLS instances	32,000	32,000
Learned MAC address entries combined for all bridge groups and VPLS instances on a chassis	64,000	64,000

Routing Protocol Maximums

Table 11 lists routing protocol maximums for the E120 router and the E320 router. The following notes are referred to in Table 11:

1. 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 three 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.
2. On the E120 router, the SRP-120 and the SRP-320 support a maximum of 64,000 IP network interfaces. On the E320 router, the SRP-320 supports a maximum of 96,000 IP network interfaces. The SRP-100 supports a maximum of 64,000 IP network interfaces.

   You can use either all dynamic interfaces or a combination of dynamic and static interfaces to achieve this maximum.
3. These values are subject to limitations on available SRP module memory, which varies according to your router configuration.
4. Depending on your configuration, the router may support more routing table entries or fewer routing table entries than this value. In any case, you can choose to limit the number of routes that can be added to the routing table on a per-VR or per-VRF basis by means of the **maximum routes** command.
5. The maximum number of ANCP adjacencies can be scaled over a maximum of 100 virtual routers. Fewer ANCP adjacencies can be scaled in configurations with more than 100 virtual routers.
6. On the E320 router, you can achieve 32,767 total Martini circuits only over Ethernet interfaces. For all routers, the total Martini circuits can be any combination of external inter-router circuits and internal circuits (local cross-connects).
7. There is no per-VR limit; all multicast routes can be on a single VR or present across multiple VRs.
8. The maximum number of interfaces can be achieved by any combination; for example, two streams each being replicated to 32,768 interfaces; 16,384 streams each being replicated four times; or any other combination.

Table 11: Routing Protocol Maximums

Feature	E120	E320
BFD		
Sessions per line module for ES2 4G LM	100	100
Sessions per line module for all modules other than ES2 4G LM	50	50
ECMP maximum paths to a destination		
BGP, IS-IS, MPLS, OSPF, RIP	16	16

Table 11: Routing Protocol Maximums Table continued

Feature	E120	E320
IPv4 forwarding table entries per chassis (See Note 1 on page 95.)	1,048,576	1,048,576
IP network interfaces (IPv4 and IPv6)		
Per chassis (See Note 2 on page 95.)	64,000	96,000
Per ES2 4G LM	16,000	16,000
Per ES2 10G LM	16,000	16,000
Per ES2 10G ADV LM	32,000	32,000
Per ES2 10G Uplink LM	8000	8000
IPv4 routing protocol scaling and peering densities (See Note 3 on page 95.)		
Routing table entries (See Note 4 on page 95.)	500,000	500,000
ANCP Adjacency Scaling (See Note 5 on page 95.)	5000	5000
BGP-4 peering sessions	3000	3000
BGP-4 routes (NLRI)	1,500,000	1,500,000
IP next hops (egress FECs); used to represent the IP addresses of next-hop routers on Ethernet interfaces	1,000,000	1,000,000
MPLS next hops (egress FECs) when graceful restart is not enabled	500,000	500,000
MPLS next hops (egress FECs) when graceful restart is enabled	250,000	250,000
MPLS forwarding entries when graceful restart is not enabled	64,000	64,000
MPLS forwarding entries when graceful restart is enabled	32,000	32,000
IS-IS adjacencies	150	150
IS-IS routes	20,000	20,000
MPLS LDP LSPs when graceful restart is not enabled	10,000	10,000
MPLS LDP LSPs when graceful restart is enabled	5000	5000
MPLS RSVP-TE LSPs when graceful restart is not enabled	10,000	10,000
MPLS RSVP-TE LSPs when graceful restart is enabled	5000	5000
OSPF adjacencies	1000	1000
OSPF routes	25,000	25,000
IPv6 routing table entries (See Note 3 on page 95.)	100,000	100,000
J-Flow statistics		
J-Flow–enabled VRs and VRFs, in any combination	16	16
Sampled interfaces per VR or VRF	32	32
Total sampled Interfaces per chassis	512	512

Table 11: Routing Protocol Maximums Table continued

Feature	E120	E320
Martini circuits for layer 2 services over MPLS		
Total Martini circuits per line module	16,000	16,000
Total Martini circuits per chassis (See Note 6 on page 95.)	16,000	32,767
External Martini circuits per chassis	16,000	32,767
Internal Martini circuits (local cross-connects) per chassis	16,000	32,767
Mobile IP bindings per chassis	–	96,000
Multicast routes (IPv4 and IPv6)		
Forwarding entries [(S,G) pairs] per chassis (See Note 7 on page 95.)	16,384	16,384
Outgoing interfaces per chassis (See Note 8 on page 95.)	65,536	65,536
Response Time Reporter simultaneous operations per VR	500	500
Response Time Reporter maximum tests per chassis (SRP-100 or SRP-320)	–	500
Response Time Reporter maximum tests per virtual router (SRP-100 or SRP-320)	–	100
VRRP VRIDs per line module	See *Ethernet VRRP VRIDs per line module* on page 92.	See *Ethernet VRRP VRIDs per line module* on page 92.

Policy and QoS Maximums

Table 12 lists policy and QoS maximums for the E120 router and the E320 router. The following notes are referred to in Table 12:

1. For more information about system resource requirements for nodes, queues, and shadow nodes, see *JunosE Quality of Service Configuration Guide, Chapter 15, QoS Profile Overview*. QoS is supported on all E Series line modules except for the ES2 10G Uplink LM.
2. For all line modules the maximum number of IPv4 or IPv6 or VLAN policy attachments is determined by the maximum number of interfaces multiplied by the number of attachment resources that are currently used. Attachment resources are only used when you attach the policy.

   The line modules support policy attachments based on the following considerations:

   - IPv4—Up to 2 ingress policy attachments and 1 egress policy attachment
   - Secure policy—Up to 1 ingress policy attachment and 1 egress policy attachment (ES2 10G LM and ES2 10G ADV LM only)
   - IPv6—Up to 2 ingress policy attachments and 1 egress policy attachment
   - IPv4 secure policy—The ES2 4G LM, the ES2 10G LM, and the ES2 10G ADV LM support up to 1 ingress policy attachment and 1 egress policy attachment
   - IPv6 secure policy—The ES2 4G LM supports up to 1 ingress policy attachment and 1 egress policy attachment
   - VLANs—Up to 1 ingress policy attachment and 1 egress policy attachment
3. Secure policies are not supported on the ES2 10G Uplink LM. IPv6 secure policies are not supported on the ES2 10G LM.

Table 12: Policy and QoS Maximums

Feature	E120	E320
QoS queues per line module (See Note 1 on page 98.)	128,000	128,000
QoS profiles configurable per chassis	1000	1000
QoS profile attachments per chassis	96,000	96,000
QoS profile attachments per line module		
ES2 4G LM	16,000	16,000
ES2 10G LM	16,000	16,000
ES2 10G ADV LM	32,000	32,000
QoS scheduler nodes per line module	64,000	64,000

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
QoS shapers per line module	64,000	64,000
Classification rules per policy	512	512
Policy classification (CLACL) entries per line module		
ES2 4G LM	256,000	256,000
ES2 10G LM	262,143	262,143
ES2 10G ADV LM	131, 071	131, 071
ES2 10G Uplink LM	131,071	131,071
Policy egress interface attachments per line module (See Note 2 on page 98.)		
ES2 4G LM combined IP and IPv6 interface attachments	16,383	16,383
ES2 4G LM combined ATM, GRE, L2TP (LAC only), MPLS, and VLAN interface attachments	16,383	16,383
ES2 10G LM combined IP and IPv6 interface attachments	16,383	16,383
ES2 10G LM VLAN interface attachments	16,383	16,383
ES2 10G ADV LM IP interface attachments	32,000	32,000
ES2 10G ADV LM VLAN interface attachments	32,000	32,000
ES2 10G Uplink LM combined IP and IPv6 interface attachments	16,383	16,383
ES2 10G Uplink LM VLAN interface attachments	8191	8191
Policy ingress interface attachments per line module (See Note 2 on page 98.)		
ES2 4G LM combined IP and IPv6 interface attachments	32,767	32,767
ES2 4G LM combined ATM, GRE, L2TP (LAC only), MPLS, and VLAN interface attachments	16,383	16,383
ES2 10G LM IP interface attachments	16,383	16,383
ES2 10G LM VLAN interface attachments	16,383	16,383
ES2 10G ADV LM IP interface attachments	64,000	64,000

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
ES2 10G ADV LM VLAN interface attachments	32,000	32,000
ES2 10G Uplink LM IP interface attachments	16,383	16,383
ES2 10G Uplink LM VLAN interface attachments	8191	8191
Rate limiters (egress) per line module		
ES2 4G LM	64,000	64,000
ES2 10G LM	64,000	64,000
ES2 10G ADV LM	64,000	64,000
ES2 10G Uplink LM	64,000	64,000
Rate limiters (ingress) per line module		
ES2 4G LM	64,000	64,000
ES2 10G LM	64,000	64,000
ES2 10G ADV LM	64,000	64,000
ES2 10G Uplink LM	64,000	64,000
Policy statistics blocks (egress) per line module		
ES2 4G LM	256,000	256,000
ES2 10G LM	256,000	256,000
ES2 10G ADV LM	512,000	512,000
ES2 10G Uplink LM	256,000	256,000
Policy statistics blocks (ingress) per line module		
ES2 4G LM	256,000	256,000
ES2 10G LM	256,000	256,000
ES2 10G ADV LM	512,000	512,000
ES2 10G Uplink LM	256,000	256,000
Parent groups (egress) per line module		
ES2 4G LM	49,151	49,151
ES2 10G LM (internal parent groups only)	8191	8191
ES2 10G ADV LM (internal parent groups only	8191	8191
ES2 10G Uplink LM (internal parent groups only)	8191	8191

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
Parent groups (ingress) per line module		
ES2 4G LM	49,151	49,151
ES2 10G LM (internal parent groups only)	8191	8191
ES2 10G ADV LM (internal parent groups only	8191	8191
ES2 10G Uplink LM (internal parent groups only)	8191	8191
Software lookup blocks per line module		
ES2 4G LM	16,383	16,383
ES2 10G LM	16,383	16,383
ES2 10G ADV LM	32,000	32,000
ES2 10G Uplink LM	16,383	16,383
Secure policies (for packet mirroring)		
Per chassis	2400	2400
Per line module (See Note 3 on page 98.)	1022	1022

Tunneling Maximums

Table 13 lists tunneling maximums for the E120 router and the E320 router. The following notes are referred to in Table 13:

1. The ES2-S1 Service IOA supports any combination of DVMRP, GRE, and L2TP tunnels up to a maximum of 8000 tunnels; however, no more than 4000 tunnels can be DVMRP or GRE tunnels in any combination.
2. For more information about supported L2TP sessions and tunnels, see *JunosE Broadband Access Configuration Guide, Chapter 11, L2TP Overview*.

Table 13: Tunneling Maximums

Feature	E120	E320
DVMRP (IP-in-IP) tunnels per chassis	4000	4000
DVMRP (IP-in-IP) tunnels per line module with shared tunnel-server ports provisioned	4000	4000
DVMRP (IP-in-IP) tunnels per ES2-S1 Service IOA (See Note 1 on page 102.)	4000	4000
GRE tunnels per chassis	4000	4000
GRE tunnels per line module with shared tunnel-server ports provisioned	4000	4000
GRE tunnels per ES2-S1 Service IOA (See Note 1 on page 102.)	4000	4000
L2TP sessions per chassis (See Note 2 on page 102.)	60,000	60,000
L2TP sessions per line module with shared tunnel-server ports provisioned (See Note 2 on page 102.)	8000	8000
L2TP sessions per ES2-S1 Service IOA (See Note 2 on page 102.)	16,000	16,000
L2TP tunnels per chassis for SRP-100	16,000	16,000
L2TP tunnels per chassis for SRP-320 with ES2 4G LM	32,000	32,000

Table 13: Tunneling Maximums Table continued

Feature	E120	E320
L2TP tunnels per line module with shared tunnel-server ports provisioned (See Note 2 on page 102.)	8000	8000
L2TP tunnels per ES2-S1 Service IOA (See Note 1 and Note 2 on page 102.)	16,000	16,000

Subscriber Management Maximums

Table 14 lists subscriber management maximums for the E120 router and the E320 router. The following notes are referred to in Table 14:

1. DHCP relay proxy maintains a list of active DHCP clients up to a maximum of 100,000 clients per chassis for all virtual routers. DHCP relay does not maintain a list of DHCP clients.

   DHCP relay proxy is notified of DHCP client deletions and subsequently deletes the client's host routes. In contrast, DHCP relay is not notified of DHCP client deletions, so the host routes for deleted clients remain in DHCP relay until you permanently delete them with the **set dhcp relay discard-access-routes** command. A maximum of 100,000 host routes for DHCP clients can be stored for all DHCP relay and DHCP relay proxy instances (that is, for all virtual routers).

2. On the E120 router, the SRP-120 and the SRP-320 support a maximum of 64,000 interfaces.

   On the E320 router, the SRP-320 supports a maximum of 96,000 interfaces. The SRP-100 supports a maximum of 64,000 interfaces.

3. For DHCPv6 local server, up to 32,000 subscribers and clients are supported on PPP/ATM and PPPoE/ATM with dynamic interfaces. Interface flapping tests have been qualified for 8000 subscribers and interfaces.

Table 14: Subscriber Management Maximums

Feature	E120	E320
DHCP external server clients (per chassis for all virtual routers; and per virtual router) (See Note 1 on page 104.)	100,000	100,000
DHCP local server (See Note 2 on page 104.)		
Client bindings per chassis	96,000	96,000
Client interfaces per chassis	64,000	96,000
Local address pools per virtual router	4000	4000
IP addresses per local address pool	96,000	96,000
DHCPv6 local server		
Clients (See Note 3 on page 104.)	32,000	32,000
DHCP relay and relay proxy client (See Notes 1 and 2 on page 104.)		
DHCP client host routes for DHCP relay and DHCP relay proxy combined (per chassis for all virtual routers; and per virtual router)	100,000	100,000
DHCP relay proxy clients (per chassis for all virtual routers; and per virtual router)	100,000	100,000
Interfaces (per chassis for all virtual routers; and per virtual router)	64,000	96,000

Table 14: Subscriber Management Maximums Table continued

Feature	E120	E320
RADIUS requests		
Concurrent RADIUS authentication requests	32,000	32,000
Concurrent RADIUS accounting requests	32,000	96,000
RADIUS route-download server downloaded routes per chassis	64,000	96,000
Service Manager		
Service definitions	2048	2048
Service sessions (active)	196,608	262,144
Active subscriber sessions	64,000	96,000
SRC Software and SDX Software		
COPS client instances	200	200
SRC clients	200	200
SRC interfaces	48,000	96,000
Subscriber interfaces (See Note 2 on page 104.)		
Dynamic subscriber interfaces per chassis	64,000	96,000
Dynamic subscriber interfaces per ES2 4G LM	16,000	16,000
Dynamic subscriber interfaces per ES2 10G LM	16,000	16,000
Dynamic subscriber interfaces per ES2 10G ADV LM	32,000	32,000
Static subscriber interfaces per chassis	64,000	96,000
Static subscriber interfaces per ES2 4G LM	16,000	16,000
Static subscriber interfaces per ES2 10G LM	16,000	16,000
Static subscriber interfaces per ES2 10G ADV LM	32,000	32,000

Appendix B

Selecting the Designated VE Device for a Multihomed Site

The information provided in this appendix updates the *BGP Multihoming for VPLS* section in *Chapter 9, VPLS Overview* of the *BGP and MPLS Configuration Guide*.

BGP Multihoming for VPLS

BGP multihoming enables you to connect a customer site to two or more PE routers to provide redundant connectivity while preventing the formation of layer 2 loops in the service provider's network. The redundant connectivity maintains the VPLS service and traffic forwarding to and from the multihomed site in the event of a PE router-to-CE device link failure, the failure of a PE router, or an MPLS reachability failure between the local PE router and a remote PE router. A redundant PE router can begin providing service to the customer site as soon as the failure is detected. BGP multihoming is very similar for both VPLS and VPWS, with only minor differences in behavior between the two L2VPN types.

When a CE device connects to multiple PE routers, each of these routers advertises reachability for the multihomed site-routes that have the same site ID in the layer2 NLRI. The other PE routers in the network use a combination of a BGP and VPLS path selection algorithm to select only one of the advertising routers to which they send traffic destined for the CE device. This path selection process eliminates layer 2 loops in the VPLS network.

Best Practice: To prevent the creation of layer 2 loops due to a misconfiguration or temporary loops during a topology change and subsequent convergence, we recommend that you employ the Spanning Tree Protocol (STP) on your CE devices.

You specify on each PE router connected to the CE device in the VPLS that the site is multihomed and you configure a priority. The priority serves as a site preference and is propagated by BGP in the local-preference attribute and in the site-preference attribute carried within the Layer2-Extended-Community.

You configure the same site ID (sometimes referred to as a VE ID) on these connected PE routers. Each of these routers then advertises reachability for the multihomed site; the VPLS NLRI contains the site ID. The site ID shared by the connected PE routers should be different than the site IDs configured on the remote PE routers in the VPLS network; if the site ID is not different, then the pseudowire will be in a site collision state. The remote routers then use the site ID to identify where to forward traffic destined for the customer site.

Although the site ID is the same for all connected PE routers, the block offset, label range, and route distinguisher can be different for each PE router. The BGP path selection process uses the route distinguisher and block offset to determine whether a layer 2 advertisement is relevant to the multihomed site. In case of unique route distinguisher per multihomed PE, VPLS NLRIs are not considered equivalent from the BGP standpoint. In consequence, a secondary path selection mechanism selects a unique path to the multihomed site. This secondary algorithm is called as the VPLS path selection process.

The PE routers run the BGP path selection process on the locally originated and received layer 2 route advertisements to establish that the routes are suitable for advertisements to other peers, such as route reflectors. For this selection process, the routes advertise different prefixes, distinguished by the site ID, block offset, and route distinguisher.

The PE router that originates the elected advertisement then becomes the designated VE device for the multihomed customer site. When the designated VE device is determined for both the local and remote customer sites for the VPLS, then a VPLS pseudowire is created between the designated VE devices.

When the PE router receives a layer 2 BGP advertisement that has the down bit set, inbound policy sets the local preference attribute to zero. The selection process can then choose an existing route from an alternate PE router, if available.

When a PE router in a VPLS domain is also a BGP route reflector (RR), the path selection process to determine the VE device for the multihomed site has no effect on the path selection process performed by this PE router for the purpose of reflecting layer 2 routes.

Layer 2 prefixes that have different route distinguishers are considered to have different NLRI for route reflection. This result of the standard BGP path selection process enables the RR to reflect all routes that have different route distinguishers to all other RR clients even though only one of these routes is used to trigger the VPLS pseudowire to the multihomed site.

Designated VE Device Selection for a Multihomed Site

BGP on each PE router in the VPLS network determines the best path to the multihomed site by comparing path attributes. The PE routers receiving the advertised routes first run the BGP selection process. The routes from the connected multihomed PE routers all share the same site ID, but can have different route distinguishers and block offsets; the routers are advertising different prefixes. The following sequences are applied to all routes on a per-prefix basis.

The designated VE device-selection process is divided into two procedures. The BGP designated VE device-selection procedure runs prior to the VPLS designated VE device-selection procedure. However, the BGP designated VE device-selection procedure takes place only when a multihomed site is being advertised from remote PEs with identical route distinguishers. If there are unique route distinguishers per PE, advertisements are not considered relevant to the BGP designated VE device-selection process. In consequence, only the VPLS designated VE device-selection process takes place.

The BGP designated VE device-selection procedure is as follows:

1. If the advertisement contains *'D' bit = 0*, then discard the advertisement.

2. Select a path with a higher preference. The preference attribute is obtained from the *site-preference*. If the *site-preference = 0*, then the preference attribute is obtained from the *local-preference*.
3. If preference values are the same, select the path with a lower router-id.
4. If router-ids are the same, routes are from the same PE. The current advertisement is considered an update.

Once the BGP designated VE device-selection procedure concludes, the VPLS designated VE device-selection procedure commences. The following steps are carried out regardless of the outcome of the BGP designated VE device-selection process:

1. If the advertisement contains *'D' bit = 0*, then discard the advertisement.
2. Select a path with a higher preference. The preference attribute is obtained from the *site-preference*. If the *site-preference = 0*, then the preference attribute is obtained from the *local-preference*.
3. If preference values are the same, select the path with a lower router-id.
4. If router-ids are the same, select the path with a lower route distinguisher.
5. If route distinguishers are the same, select the path with a lower block offset value.
6. If block offset values are the same, the current advertisement is considered an update.

The result of this process establishes that the best path is suitable for establishing a pseudowire from the remote PE router to the PE router. That PE router is accordingly selected as the designated VE device.

When the remote PE router establishes or refreshes a pseudowire to the local PE router, it verifies whether the prefix is in the range required for the site ID based on the block offset and label range advertised by the designated VE device. If the prefix is out of range, then the pseudowire status is set to OR (out of range).

One of the following cases applies for each PE router when it completes the BGP path selection process for a layer 2 advertisement on the VPLS.

- The PE router originated one of the multihomed advertisements and selected its own advertisement as the best path.

  This PE router hosts the designated VE device. Selection as the designated VE device triggers the creation of pseudowires to and from the other PE routers in the VPLS. When the remote customer site is also multihomed, then the designated VE device triggers the creation of pseudowires to and from only the designated VE device for the remote site.

- The PE router originated one of the multihomed advertisements but did not select its own advertisement as the best path.

  This PE router is one of the redundant PE routers for the multihomed site; it does not host the designated VE device. If its status has just transitioned from being the designated VE device, then the PE router tears down all the pseudowires that it had to and from the other PE routers in the VPLS network.

- The PE router receives the multihomed advertisements and selects a best path; it does not originate any of these advertisements because it is not connected to the multihomed customer site.

If the selected best path-and therefore the designated VE-has not changed, then nothing happens. If the best path has changed, then this PE router brings up pseudowires to and from the new designated VE device and tears down the pseudowires to and from the previous designated VE device.

If this PE router does not select a best path after running the process, then the local PE router does not consider the remote site to exist.

When a VE device receives an advertisement for a layer 2 NLRI that matches its own site ID but the site is not multihomed, then the pseudowire between it and the transmitting PE router transitions to a site collision (SC) state and is not considered to be up.

Multihoming Reaction to Failures in the Network

The redundant connectivity provided by a multihoming configuration protects against several types of network failure.

- CE-Link failure between the CE device and the PE router-BGP on the PE router is notified when the circuit goes down. BGP then modifies the circuit status vector bit in the MP_REACH_NLRI to indicate that the circuit is down.

  If all VPLS local attachment circuits are down, then BGP modifies the down bit in the VPLS advertisement Layer2-Extended-Community to state that the site is down. When the bit is modified, BGP advertises the route to all remote PE routers to inform them that the circuit (and site) is down. The remote PE routers each run the best path selection process again and adjust the VPLS pseudowires as needed.

- Failure of MPLS reachability to the remote PE router-BGP on the PE router is notified that MPLS connectivity to the BGP next hop is gone. BGP then modifies the circuit status vector bit in the MP_REACH_NLRI to indicate that the LSP is down. When the bit is modified, BGP advertises the route to all remote PE routers to inform them that connectivity is down from the local site to the remote site.

  The down bit is set if no remote PE router is reachable by MPLS. This enables the remote PE routers to consider the other multihomed PE router as the designated VE device for the multihomed-site.

  The remote PE routers each run the best path selection process again and adjust the VPLS pseudowires as needed.

- PE router failure-When either the PE router or its BGP process fails, peer PE routers detect expiration of the holdtimer and bring down their peering sessions, and remove layer 2 advertisements from the PE router. Alternatively, the PE routers can detect unreachability to the BGP next hop that represents the failed PE router. In this case the peer routers mark the layer 2 routes advertised by PE router as unreachable. The peer PE routers each run the best path selection process again and adjust the VPLS pseudowires as needed.

  A similar response results when you adjust the multihoming priority of the PE routers connected to the multihomed site, effectively performing and administrative failover to another PE router. BGP sends a layer 2 update with the new local preference attribute to all peer PE routers. The peer PE routers each run the best path selection process again and adjust the VPLS pseudowires as needed.

To modify their pseudowires, the peer routers correct their MPLS forwarding tables and set up new entries in their pseudowire tables.