



JunosE™ Software for E Series™ Broadband Services Routers

Release Notes

Release

11.3.3



Published: 2011-12-29

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

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

Products made or sold by Juniper Networks (including the ERX310, ERX705, ERX710, ERX1410, ERX1440, M5, M7i, M10, M10i, M20, M40, M40e, M160, M320, and T320 routers, T640 routing node, and the Junos, JunosE, and SDX-300 software) or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Copyright © 2011, Juniper Networks, Inc.

All rights reserved. Printed in USA.

JunosE™ Software for E Series™ Broadband Services Routers Release Notes, Release 11.3.3

Revision History

December 2011—FRS JunosE 11.3.3

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

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

Release 11.3.3	1
Release Installation	1
Upgrading to Release 5.3.0 or a Higher-Numbered Release	1
Upgrading from Release 5.1.1 or Lower-Numbered Releases to Release 6.x.x or Higher-Numbered Releases	1
Moving Line Modules Between Releases	2
SRP Module Memory Requirements	2
Hardware and Software Compatibility	3
Requesting Technical Support	3
Self-Help Online Tools and Resources	3
Opening a Case with JTAC	4
Release Overview	5
Before You Start	5
Release Highlights	6
Early Field Trial Features	6
DHCP	6
Stateful Line Module Switchover (High Availability)	7
Unsupported Features	8
E120 Router and E320 Router	8
Policy Management	8
Stateful SRP Switchover (High Availability)	9
Release Software Protocols	9
Core Routing Stack	9
Layer 2 Protocols	9
Multiprotocol Label Switching (MPLS)	9
Network Management Protocols	9
Routing Protocols	10
Security Protocols	10
SRC Software and SDX Software Compatibility Matrix	10
Known Behavior	11
AAA	11
ATM	11
BGP	11
BGP/MPLS VPNs	12
B-RAS	12
CLI	12
DHCP	16
DHCP External Server	16
Ethernet	17

Flash	18
GRE	18
Hardware	19
HDLC	19
IP	20
IPSec	21
IS-IS	22
L2TP	23
Line Module Redundancy	23
MLPPP	24
MPLS	24
Multicast	24
Packet Mirroring	25
Policy Management	26
PPP	28
PPPoE	28
QoS	28
RADIUS	29
SNMP	29
SRC Software and SDX Software	30
SSH	30
Stateful SRP Switchover (High Availability)	31
System	32
Tunneling	33
Known Problems and Limitations	33
ATM	33
BFD	33
DHCP	33
DHCP External Server	34
DoS Protection	34
Forwarding	34
GRE	36
ICR	36
IGMP	36
IP	37
IS-IS	37
L2TP	37
LDP	37
MLD	38
MPLS	38
Multicast	38
Policy Management	38
PPPoE	40
QoS	40
RSVP-TE	41
Service Manager	41
SRC Software and SDX Software	41
Stateful SRP Switchover (High Availability) and IP Tunnels	41
Subscriber Management	42
System	43

TCP	43
Unified ISSU	44
Resolved Known Problems	44
ATM	44
BGP	44
DHCP	44
DHCP External Server	44
Dynamic Connection Manager	45
Ethernet	45
Forwarding	45
Hardware	46
HTTP	46
IGMP	46
IP	46
IPv6	47
IS-IS	47
L2TP	47
LDP	47
MPLS	47
OSPF	48
Policy Management	48
PPP	48
PPPoE	49
QoS	49
RSVP-TE	49
SNMP	49
SRC Software and SDX Software	49
SSH	50
System	50
TCP	50
Unified ISSU	51
Errata	51
Appendix A System Maximums	69
ERX310, ERX7xx, and ERX14xx System Maximums	70
General System Maximums	70
Physical and Logical Density Maximums	71
Link Layer Maximums	74
Routing Protocol Maximums	79
Policy and QoS Maximums	82
Tunneling Maximums	85
Subscriber Management Maximums	87
E120 and E320 System Maximums	90
General System Maximums	90
Physical and Logical Density Maximums	91
Link Layer Maximums	93
Routing Protocol Maximums	98
Policy and QoS Maximums	101
Tunneling Maximums	105
Subscriber Management Maximums	107

Release 11.3.3

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 11.3.3 of the system software for the Juniper Networks E Series Broadband Services Routers and contain the following sections:

- *Release Highlights* on page 6
- *Early Field Trial Features* on page 6
- *Unsupported Features* on page 8
- *Release Software Protocols* on page 9
- *SRC Software and SDX Software Compatibility Matrix* on page 10
- *Known Behavior* on page 11
- *Known Problems and Limitations* on page 33
- *Resolved Known Problems* on page 44
- *Errata* on page 51
- *Appendix A, System Maximums*, on page 69

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 11.3.2 and 11.3.3. Before you use your new software, read these *Release Notes* in their entirety, especially the section *Known Problems and Limitations*. You need the following documentation to fully understand all the features available in Release 11.3.3:

- These 11.3.3 *Release Notes*, which describe changes between Release 11.3.2 and Release 11.3.3
- The 11.3.2 *Release Notes*, which describe features available in Release 11.3.2
- The 11.3.x documentation set, which provides detailed information about features available in Release 11.3.0

The 11.3.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	<i>ERX Hardware Guide</i> <i>E120 and E320 Hardware Guide</i>
Learn about modules	<i>ERX Module Guide</i> <i>E120 and E320 Module Guide</i> <i>E Series End-of-Life Module Guide</i>
Get up and running quickly	<i>E Series Installation Quick Start poster or ERX Quick Start Guide</i> <i>E120 and E320 Quick Start Guide</i>
Configure the router	<i>JunosE System Basics Configuration Guide</i>
Configure physical layer interfaces	<i>JunosE Physical Layer Configuration Guide</i>
Configure link layer interfaces	<i>JunosE Link Layer Configuration Guide</i>

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

Release Highlights

Release 11.3.3 is a maintenance release.

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
```

Stateful Line Module Switchover (High Availability)

- Support for Stateful Line Module Switchover on ES2 4G LMs with ES2-ES1 Service IOA on E120 and E320 Routers

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

Stateful line module switchover reduces the impact on subscriber traffic during a stateful switchover from the active line module to the standby line module by ensuring that existing subscriber sessions remain active with a brief disconnection in traffic. Stateful line module switchover maintains user sessions and reduces data forwarding outage to a brief duration through the router during the switchover, thereby improving the overall availability of the router.

The following commands have been added or modified to support configuration and monitoring of stateful line module switchover:

- **line-card switch**
- **mode high-availability slot**
- **show redundancy history line-card**
- **show redundancy line-card**

As part of this feature, the following SNMP MIB objects have been added to the Juniper Networks Redundancy MIB to support the stateful line module switchover functionality:

- junilcRedundancyActiveSlot
- junilcCardRedundancyStandbySlot
- junilcRedundancyLastResetReason
- junilcRedundancyActivationType
- junilcRedundancyHaActiveTime
- junilcRedundancySwitchoverTime

In addition, the ha event log category has been modified to record information when the system transitions from one high availability state to another, when a stateful line module switchover is performed.

Unsupported Features

The JunosE Release 11.3.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
4.0.0R3	Not applicable	10.3, 11.0, 11.1
4.0.0R7	Not applicable	10.3.3, 11.3.1, 12.0.0, 12.1.1
4.1.0	Not applicable	12.0.1, 12.1.1, 12.2.0

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.

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

For a sample configuration scenario that illustrates how to configure hold timers for successful graceful restart in a scaled environment, see *JunosE BGP and MPLS Configuration Guide, Chapter 1, Configuring BGP Routing*.

- 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.

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	
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 crt	Global Configuration	[no] ipsec crt
interface hssi	Global Configuration	
invert tx clock	Global Configuration	

Deprecated Command	Command Mode	Preferred Command
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	
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

Deprecated Command	Command Mode	Preferred Command
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
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

- 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 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.

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 a 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.

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

- 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]

- 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 longname
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.
- 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.
- 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.
- 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.

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]

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 restarting 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. [Defect ID 90546]

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.

L2TP

- 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.
- 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.

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.

Multicast

- 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.
- 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.

Packet Mirroring

- 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.
- 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 policies that were attached by a trigger action. If the interface-specific secure policies 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.

Policy Management

- The ES2 10G LM does not support the deprecated **next-hop** command.
- 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 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*.
- 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.

- Although it is not required, you can enclose the name of the classifier when you use the **show classifier-list** *classifierName* command and the name of the policy list when you use the **show policy-list** *policyName* command within double quotation marks. This method of specification of policy and classifier names ensures that the CLI interface does not process the abbreviated forms of the names as system-defined keywords, such as **brief** and **detailed**, available with **show policy-list** and **show classifier-list** commands.

For example, if you specify the **show policy-list b** command without enclosing the letter "b" within double quotation marks, assuming a policy list with the name "b" has been configured, the system auto-completes the letter "b" as **brief** and considers the command to denote a condensed display of policy lists (equivalent of **show policy-list brief** command). Similarly, if you enter the **show classifier-list d** command to display the details of a configured classifier list with the name "d", the CLI interface processes the command as a listing of classifier details (equivalent of **show classifier-list detailed** command).

To avoid incorrect and unexpected behavior in the output of the **show classifier-list** *classifierName* and **show policy-list** *policyName* commands, you must enclose the names of policy lists and classifier lists while using these commands within double quotation marks, especially if the names of the policy and classifier lists begin with letters that match the auto-complete forms of keywords. If the names of the policy and classifier lists do not match the beginning letters of the keywords or if you enter the full names of the policy and classifier lists, the system accurately processes the names even if you do not enclose them within double quotation marks while using these commands.

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.

- In Release 7.2.0 and higher-numbered releases, you can configure the simple shared shaper to select scheduler nodes in a named traffic-class group as active constituents.

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.

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
juniApsEventSwitchover	.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
juniApsEventModeMismatch	.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
juniApsEventChannelMismatch	.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
juniApsEventPSBF	.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
juniApsEventFEPLF	.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.

SRC Software and SDX Software

- In a network in which approximately 40,000-45,000 IP interfaces are managed by an SRC client on an E Series router, if you enter the **sscc enable** command to enable the SRC client after it was previously disabled, the CLI interface stops responding and is not accessible for about 15 minutes. [Defect ID 187946]

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

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.

Tunneling

- When you configure the GE-2 line module or the GE-HDE line module with a shared tunnel-server port, the available bandwidth for tunnel services is limited to 0.5 Gbps per module. When you configure the ES2 4G line module with a shared tunnel-server port, the available bandwidth for tunnel services is limited to 0.8 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. Problems that have not been reported by customers are documented only in these Release Notes.

ATM

- 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.

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]

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

- 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 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.

Forwarding

- After you configure fast reroute extensions to RSVP-TE to enable local protection for the ingress router of the primary LSP by using bypass tunnels, forwarding of IPv6 traffic to some of the labeled BGP routes or IPv6 destinations over ES2 10G LMs and ES2 4G LMs fails. [Defect ID 189451]
- An ES2 10G LM resets when you attempt to remove a multicast port in a dual-stack environment that contains both IPv4 and IPv6 subscribers. [Defect ID 189867]
- On an E120 or E320 router that is running JunosE Release 11.2.1, the ES2 10G LM with an ES2-S2 10 GE PR IOA stops egress forwarding because of an egress IXP static RAM (SRAM) parity error. This problem is observed during the removal of the LM from its slot and reinsertion into the same slot when subscriber sessions are established. This behavior also happens when about 6000 subscribers are brought up and user traffic is transmitted. This problem depends on the type of user traffic sent. [Defect ID 189846]
- 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 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]

- 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.

- 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]
- L2TP subscriber sessions gradually terminate when the following conditions are true: [Defect ID 188850]
 - MPLS is configured using equal-cost multipath (ECMP) links between a LAC device and an LNS device in an L2TP tunnel.
 - The LNS device contains an ES2 10G ADV LM with Service IOA as the tunnel server module.
- 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
```

- If you perform unified ISSU from JunosE Release 11.3.0 to a higher-numbered release on a router chassis with an ES2 10G ADV LM that contains a shared tunnel-server port or a dedicated tunnel-server port, the LM resets. The existing dedicated and shared tunnel-server sessions that were configured to terminate on the ES2 10G ADV LM are dropped when the LM resets. The LM comes back online, running the release to which it was upgraded, after it completes the generation of a core dump file. Also, new tunnel-server sessions are established correctly after unified ISSU is completed. [Defect ID 187728]

GRE

- After a stateful SRP switchover operation, if you delete the shared tunnel-server port configured on an LM and configure the shared tunnel-server port on another LM, the GRE tunnels are not transferred from the LM on which the tunnel-server port was deleted to the LM on which the tunnel server was configured. The removed tunnels do not come up until you perform an event that triggers the restoration of the GRE tunnels, such as addition of new tunnels, deletion of other existing tunnels, or configuration of an LM on the router chassis. This problem occurs only with GRE tunnels after the completion of a stateful SRP switchover procedure. [Defect ID 187804]

Work-around: To correct this problem, create a dummy GRE tunnel interface for temporary purposes using the **interface tunnel** command. Creation of the dummy tunnel interface causes the previously-configured GRE tunnel-service interfaces to be restored on the changed shared tunnel-server port on the LM. You can delete the dummy tunnel interface after all the GRE tunnel interfaces are up.

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.

IGMP

- 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.)

- 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]
- IGMPv3 proxy is not supported. [Defect ID 46038]

IP

- 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]

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

- L2TP subscriber sessions that were previously established are disconnected when you perform a stateful SRP switchover operation in a scaled environment. [Defect ID 187454]
- When you perform a stateful SRP switchover procedure on an LNS device that contains an ES2 4G LM with Service IOA (tunnel server module), some of the 16,000 subscriber sessions that are active over 16,000 tunnels are terminated. This problem occurs when OSPF is used as the routing protocol between the LAC and LNS devices in the L2TP tunnel, and with the number of L2TP retransmission attempts configured as 10. [Defect ID 187358]
- 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]

LDP

- Some of the LDP sessions on ATM line modules do not come up when you perform a unified ISSU operation on an ERX router, which is the restarting router, from JunosE Release 10.3.3 to Release 11.3.2. This problem occurs when an ERX router functions as the restarting router and an E120 or E320 router functions as the helper router. [Defect ID 189400, 189588]

MLD

- MLDv2 proxy is not supported. [Defect ID 46038]
- 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 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.)

MPLS

- 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]
- In a scaled environment with a large number of MPLS RSVP-TE tunnels configured, the states of the hello adjacency instances in the State field in the output of the `show mpls rsvp hello instance` command are displayed as Down for loopback interfaces. The correct behavior is that the RSVP-TE hello adjacencies must always be in the Up state for loopback interfaces. [Defect ID 189565]
- 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`

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]
- 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.

- 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.

- When you perform a stateful SRP switchover with high availability in the enabled state and with approximately 5000 dual-stack subscriber sessions, independent IPv4 sessions, or independent IPv6 sessions established on the router, the following log message is recorded for the policyMgrGeneral system logging category: [Defect ID 186570]

```
ERROR 09/26/2011 22:30:43 policyMgrGeneral: Error restoring policy attachment for
480926 from MS/NVS
```

This problem occurs when the router configuration contains GRE tunnels, IPv4 secure policies, and IPv6 secure policies, and packet mirroring is enabled using username as the trigger. This problem might also happen during unified ISSU.

- If you modify the reference-rate policy-parameter, by using either the **policy-parameter reference-rate** command at the interface level or the global level or the RADIUS-initiated change-of-authorization (CoA) method, for rate-limit profiles that are used in internal parent groups, the modified rate-limit profile is not considered during the forwarding of traffic on ES2 10G LMs, ES2 10G Uplink LMs, and ES2 10G ADV LMs. The old, stale rate-limit profile is used for forwarding of traffic. This problem occurs with both IPv4 and IPv6 traffic. This behavior is not observed with ES2 4G LMs. [Defect ID 185409]

In such a scenario, although the updated reference-rate policy-parameter value is reflected in the output of the **show ip interface** and **show ipv6 interface** commands, the modified reference-rate is not used for processing of forwarded traffic.

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

- 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 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]
- 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.

- 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 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.

- 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]

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.

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

- When you perform two stateful SRP switchover operations, high availability state transitions to disabled, and DHCP and PPP subscribers cannot log in to the router. This problem occurs when the following conditions are met: [Defect ID 86809]
 1. Approximately 15,000 DHCP and PPP subscribers are logged in to the router and the subscribers are managed by the SRC software.
 2. A stateful SRP switchover is performed, and after the high availability state changes to active on completion of the stateful switchover, another stateful SRP switchover is performed.
 3. After the completion of the second stateful SRP switchover, the logged-in subscriber sessions are terminated and an attempt to log in the subscribers is performed again.

This problem occurs only with SRP-10G modules and not with other SRP modules.

- 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 187444]

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 configure 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 causes 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 configure a static interface column and remove 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

- Memory leak is observed with the SRP-100 module while subscribers are being brought up on a LAC device and the active link between the LAC device and the LNS device in an L2TP tunnel is flapping. This problem occurs when the following steps are performed: [Defect ID 189353]
 1. Two redundant links connect the LAC device to the LNS device in the L2TP tunnel.
 2. DHCPv6 subscribers over PPPoE interfaces connected to a LAC device are attempted to be brought up.
 3. The active link between the LAC and LNS devices flaps continuously 1000 times using the shutdown and no shutdown commands.
 4. Memory-related output information is collected at a base condition where the active link is up again and no subscriber is connected to the router.

When you perform each iteration of the preceding four steps, the amount of free memory on the SRP-100 module decreases and validates a memory leak.

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

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

Resolved Known Problems

Release 11.3.3 is based on the 11.3.2 FRS release and incorporates all problem resolutions found in that release. For information about resolved problems in a patch release, customers with valid service agreements may log in to the JunosE Download Software page on the Customer Support Center Web site at <https://www.juniper.net/support/products/junose>. Select the Patch Release History for the JunosE release you are interested in.

The following problems were reported open in Release 11.3.2 and have been resolved in this release, or have been resolved since the 11.3.2 FRS release. For more information about problems in this list that were reported by customers, 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. Problems that have not been reported by customers are documented only in these Release Notes.

ATM

- In the output of the **show atm vp-tunnel *interfaceSpecifier*** command, the Type column always displays the label "Vbr", regardless of whether constant bit rate (CBR) or variable bit rate (VBR) is configured as the service category on an ATM PVC. However, the service category of an ATM PVC is correctly displayed in the output of the **show configuration interface atm** command. [Defect ID 90805]
- OC3-4A line module reset:reset type: panic; file: tm1483Dynamic.cc; line: 2466; task: scheduler. [Defect ID 91879]

BGP

- High SRP CPU utilization because of BGP. [Defect ID 91868]
- SRP 320 reset: reset type: panic; task: bgp; file: bgpPrefix.cc. This problem occurs when a router is configured with BGP signaling for VPLS and the BGP peer sends updates for the L2VPN prefix with a block offset set to 0. [Defect ID 91706]
- BGP connections flap due to the TCP error code of 32 (Broken pipe), which can be seen from the output of the **show ip bgp neighbor** command and also viewed in the bgpConnections and bgpMessages system logging category events. [Defect ID 92086]

DHCP

- Subscribers are not able to connect when the DHCP proxy client stops DORA (transmission of discovery, offer, request, and acknowledgment messages between the server and the client) process due to stale FSM/CSM entries [Defect ID 92037]

DHCP External Server

- The dynamic subscriber interface (DSI) column might be deleted after the client sends a rebind request with a different xid. [Defect ID 91484]
- DHCP external bindings are deleted after multiple renew requests are received from the clients with different xids in quick succession. [Defect ID 92573]

Dynamic Connection Manager

- Service manager reports deactivation or activation fails on every other service request for redirected HTTP sessions. [Defect ID 91749]

Ethernet

- PPPoE subscribers that are connected over a LAG bundle are under the Ethernet member links, instead of being displayed against the LAG bundle when the **show subscribers** command is entered. This problem does not occur for DHCP subscribers logged in over a LAG bundle. The routing table and other PPPoE-related statistics display the subscribers correctly as being logged in over the LAG bundle. [Defect ID 91940]

Forwarding

- Although the customer premises equipment (CPE) or the sender adds variable-length information in the Data field in the PPP echo request that it transmits to the BRAS server, the server uses a different value in the echo reply it sends to the CPE. This behavior causes instability for certain CPEs. [Defect ID 93014]
- ARP does not work correctly for member links in the LAG bundle that are configured as redundant interfaces when they encounter a failure and fail over to other links in the LAG bundle. [Defect ID 91534]
- OC3 Line Module reset type: processor exception 0x300 (data access: address not mapped (read attempt)); task: Ingress FC[4]. This reset of the OC3 LM occurs in a BRAS environment. [Defect ID 92766]
- ARP requests that are received on a VLAN subinterface in the administratively down state are forwarded to the interface controller for being processed. [Defect ID 90332]
- When IPv4 and IPv6 addresses are configured on two different virtual routers, removal of the IPv6 address on a 10-Gigabit Ethernet interface on an ES2 10G Uplink LM causes the ARP application to not work properly and traffic forwarding to be disrupted. This problem does not occur with Gigabit Ethernet interfaces. [Defect ID 91431]
- GE/FE line modules reset due to ill-formed or oversized packets generated by the module facing Link ram issue-IFA fatal panic history: 43008 [Defect ID 91785]
- L2TP mirror does not have the correct packet offset for transmitted data on ES2 10G LMs for L2TP subscribers. [Defect ID 91803]
- ES2 10G ADV LM reset; file: ic1Detector.cc; camClassifierInterfaceIngress: 1; Interface Bus Parity Error. [Defect ID 91864]
- ES2 10G ADV LM resets incorrectly when a static RAM (SRAM) parity error occurs. [Defect ID 91907]
- DHCPv6 clients cannot be logged in because the solicit messages from the clients are dropped and not processed by the delegating router if the subscriber interface is configured to demultiplex the source address. This problem does not occur if the demultiplexing method is not based on source address. [Defect ID 92693]
- Memory leak is observed on ES2 10G LMs because of memory allocation failure for "luInstructions:nextHop fifty six byte objects" (where the total allocated memory has been utilized) while building the egress branch for PPP interface. [Defect ID 92731]

Hardware

- Redundant IOA on an ES2 4G LM displays the question mark “?” character in the output of the **show hardware** command when stateful SRP switchover operation takes place. There is no impact on the regular router processing. When the line module takes over any slot in the redundancy group, it displays the correct hardware. [Defect ID 92418]
- ES2 10G ADV LM reset; file: ic1Detector.cc; camClassifierInterfaceIngress: 1; sraPanics: 3072; ingressIxpMePanics: 134217728 [Defect ID 91863]
- Add the option to take line cards experiencing link RAM errors off line automatically. [Defect ID 91890]
- ES2 10G ADV LM processor exception 0x300(data access: address not mapped (read attempt)) [Defect ID 91981]
- In the output of the **show interface gigabitEthernet** command on ES2 10G LMs, incorrect statistics are displayed. Although the MAU field displays that a port is not being used and is empty, the data rates of traffic received and sent display a significant number. Although the traffic statistics on ports that are in use is correctly displayed, the unused ports do not become active at all, the physical link LED is always turned off, and the peer device cannot come up. [Defect ID 92293]
- SRP reset type: panic arg (0xa400c4) file: pifUtils.cc task: dprofColl . [Defect ID 92714]

HTTP

- When the memory heap utilization on an SRP module is high (about 99% of its total memory block), the HTTP application becomes unresponsive. In such a case, when the HTTP local server on the virtual router is attempted to be removed, the "resources not available" error is displayed. [Defect ID 190427]

IGMP

- If IPv4 hardware multicast packet replication is enabled on a Gigabit Ethernet interface using the **ip multicast ioa-packet-replication gigabitEthernet** command, this setting is not displayed in the output of the **show configuration** command. [Defect ID 92572]

IP

- Only one MTU changes in an IPv6/IPv4 dual stack. IPv6 MTU cannot be adjusted automatically according to the physical MTU. [Defect ID 91245]
- System logging initialization with an unexpected UID value in the routerId.cfg file causes an indefinite system outage. [Defect ID 91528]
- GE-HDE (2 ports) LM card reset type: panic, arg (0x1a9fcb5) file: ppp.cc line: 14230 [Defect ID 91947]
- The SRP module on an E320 router resets when it receives a packet with an invalid IP-Option header. [Defect ID 92185]
- High CPU utilization is noticed on access GE line modules owing to slow processing of ARP traffic. This problem occurs with DHCP external server, DHCP relay, and dynamic subscriber interfaces configured with GE LMs used for terminating DHCP subscribers. [Defect ID 92539]

IPv6

- When static routes with null 0 interfaces as the next hop addresses are configured, the SNMP walk of inetCidrRouteTable enters into a loop. [Defect ID 187366]
- IPv6 stops traffic forwarding after reconfiguring the IPv6 interface addresses. [Defect ID 91490]
- In a dual-stack environment in which both IPv6 and IPv4 subscribers are present, the call setup rate for subscribers logging in using PPPoE links is very low, approximately 10 calls per second. In a sample scenario in which 48,000 dual-stack subscribers using PPPoE connections are attempted to be logged in, about 1 hour and 20 minutes is required to log in the clients. This problem occurs when IPv6 addresses are assigned to subscribers using Neighbor Discovery router advertisements and Prefix Delegation method is not used. [Defect ID 92390]
- An IPv6 address cannot be added using the **ipv6 address** command in Interface Configuration mode with an even number in the least significant bit (LSB) of a /127 subnet mask. [Defect ID 92506]

IS-IS

- In MPLS VPN environment, the **slot erase** command does not work properly when the configuration on the Gigabit Ethernet VLAN subinterface on the slot includes IS-IS and MPLS protocols. The command fails and the Telnet session from which the command is entered stops responding. [Defect ID 91295]
- Adding an unnumbered interface on IS-IS instance as a passive interface causes high CPU utilization of about 100 percent. [Defect ID 92468]

L2TP

- L2TP LAC: Hello control message with AVP reserved bits all set is responded by a wrong StopCCN Error code. [Defect ID 91289]
- L2TP tunnel failover on a LAC device takes approximately 15 minutes when a remote link on the LNS server fails. [Defect ID 92451]

LDP

- If MPLS LDP FEC deaggregation to bind a separate label to each prefix on a virtual router is configured by using the **mpls ldp deaggregate** command, the LDP egress router advertises the non-null label and not the implicit null label, by default, to the upstream neighbor. [Defect ID 91556]
- On an E Series router that has two ES2 10G LMs as uplink LMs and is connected to two other uplink routers with OSPF, LDP and BGP peerings or adjacencies, when the downlink interface's configuration is deactivated on one of the M10 routers, the LDP hellos and keepalives sent from the E Series router are delayed. The ERX router also does not process the hellos sent by the M10 router correctly and wrongly terminates certain sessions with a logged notification. [Defect ID 91268]

MPLS

- ES2 10G Uplink LM reset type: processor exception 0x68616c74 (halt) task: scheduler [Defect ID 91482]

OSPF

- If the OSPF opaque LSA contains the Router Address TLV, the maximum link-stage age displayed under the Age field in the output of the **show ip ospf database** command is always 3600 seconds and this value is not refreshed properly to show the correct age. [Defect ID 91496]
- An RSVP-TE tunnel might remain in the down state if a router LSA signals a metric of 65535 for one of the neighbors that is being considered during the processing of the Constrained Shortest path First (CSPF) algorithm. [Defect ID 91579]
- An E Series router that functions as an area border router (ABR) experiences high CPU utilization constantly for the OSPF application when LSAs are generated by the ABR. [Defect ID 92143]

Policy Management

- Although a carriage return, <cr>, is displayed when a question mark (?) is typed after entering the **vlan classifier list classifierName** command without defining any other keyword or CLACL option, an error message is displayed stating that a VLAN classifier list cannot be configured without any classification criteria when **Enter** is entered to configure the VLAN CLACL with only the name. [Defect ID 184139]
- The "Ping/ICC monitoring threshold exceeded" message is generated and the line module undergoes a reset. This problem occurs with DHCP subscribers when the classifier list used in a policy list attached to interfaces on which subscribers are logged in is modified. [Defect ID 191289]
- After a stateful SRP switchover on a router acting as an SRC client, provisioning of policies from the COPS server to the SRC client fails. The default output policy configured on the COPS server is not properly applied to users logged in before the stateful switchover. Also, the defined service policies on the COPS server are not applied to new users logging in after the default output policy is configured. [Defect ID 92919]
- If a secure policy is already attached to subinterfaces when a DHCP subscriber logs in, the junipacketMirrorSessionStart SNMP notification is not triggered. [Defect ID 92286]
- RADIUS request messages for ingress and egress policy attachment fails on IPv4 interfaces when a subscriber sends a dual-stack request for both IPv4 and IPv6 interfaces. RADIUS requests work properly when ingress and egress policies are attached to IPv4 interfaces for subscribers that are connected only using IPv4. [Defect ID 91506]

PPP

- PPP fragmentation and PPP reassembly settings that were previously configured on the router are deleted after an upgrade operation. [Defect ID 189802]
- In a BRAS environment, if the PPP MRU configured in a profile is 1500 bytes and the LCP configuration request AVP from the LAC is sent with an MRU of 1492 bytes, the PPP LCP negotiation of the MRU encounters problems. [Defect ID 91772]
- The name of a virtual router (VR) configured as the authentication VR context using the **ppp authentication virtual-router** command is incorrectly displayed with a mismatching VR name in the PPP profile configuration section in the output of the **show configuration** and **show profile** commands. [Defect ID 89489]

PPPoE

- PPPoE sessions not established owing to insufficient system resources after an upgrade. The PPPoE Active Discovery Initiation (PADI) packets received from the clients are not responded. This problem occurs when the setting for the maximum number of PPPoE sessions is not configured in the dynamic profile for PPPoE subscribers over ATM or by the RADIUS server. [Defect ID 92744]

QoS

- ES2 4G LM reset at file: c2QosAgentIcc.cc message: file fc2QosAgent.cc line 3173: Missing node for [if0xaec00fff TCG0] at lvl 1 [Defect ID 91889]

RSVP-TE

- If an MPLS minor interface is attempted to be deleted using the **no interface tunnel mpls** command, RSVP-TE does not remove the interface and it is still displayed as configured in the output of the **show mpls minor-interface** command. [Defect ID 92469]
- A separate command with the syntax, **show mpls rsvp**, to display MPLS RSVP-TE status and configuration is not available. [Defect ID 92820]

SNMP

- The output of the **show snmp** command indicates a discrepancy between the counters displayed for the Get-Request PDUs and Get-Response PDUs fields when Get Bulk requests are used by the client to obtain the identifiers and values as a group rather than one value at a time. [Defect ID 92278]

SRC Software and SDX Software

- When a subscriber managed by the SRC server logs out and the COPS session is disconnected, the logged out subscriber's entry is retained in the SAE database. As a result, when any other subscriber attempts to log in and the COPS session is reestablished between the SRC client and the SRC server, access is denied and the second UserSession log is recorded. [Defect ID 92528]
- The proprietary JunosE Policy Information Base (PIB) attribute is not contained in the Enterprise SNMP Management Identifiers (JUNIPER-UNI-SMI) file. [Defect ID 90238]
- The SRC client does not send Request packets for additional client bindings after a network service interruption. The Request packets are sent only to the DHCP hosts that are in a steady state prior to the network outage and not for the hosts that are added while the network connection is down. [Defect ID 91737]
- When the hostname of an E Series router is modified to be different from the hostname that is defined on the SRC server and SRC client is running on the router, flapping of COPS sessions is noticed and eventually causes the system resources of the SRC client to be exhausted. Although the SRC client retries connection establishment with the SRC server, the SRC client does not wait until the configured retry timer value is reached to reattempt the creation of the connection. [Defect ID 92051]
- Because of many stale entries in the SscDhcpHandleTable, the DHCP local server operating in equal-access mode does not lease IP addresses to clients correctly. [Defect ID 92199]

- Because the router that functions as the SRC client does not send the Delete Request (DRQ) message to the SRC server, when the second user attempts to log in, the authentication of the second user fails. This problem occurs because the first user session that was established is still preserved in the SRC server even after that user has logged out and when the second user receives the same IP address as the first user, the SAE rejects the authentication request. [Defect ID 92283]
- SRP reset: task: sshd; file: sshServer.cc; line: 683; last errno: 0x0; lr: SshServer::Daemon::taskMain [Defect ID 92343]
- Incorrect or duplicate update policy request messages are sent from the SRC client to the COPS server. [Defect ID 92442]

SSH

- While an SSH session is attempted to be established with the router, the SSH sessions are blocked and do not get cleared after the **disconnect ssh sessionid** command is entered. The output of the **show users** command does not display information about these blocked SSH sessions. [Defect ID 92159]

System

- SRP reset at task: dprofColl Srp40 InternalNetwork :: createChannel (ArInternalNetwork::ControlSession). With router that contains Gigabit Ethernet LMs on the access side, when 48,000 DHCP clients are brought up and the allocated addresses or bindings are released after they are logged out, and the configured DHCP relay and DHCP external server are removed, the active SRP module resets when the LMs are reloaded. This problem also occurs with PPPoE and L2TP subscribers. [Defect ID 186867]
- When the Service Manager application is enabled on a router in a BRAS environment, an SRP-40 module resets when a stateful SRP switchover is performed. [Defect ID 92557]
- The hold-off timer stops during its countdown in the middle and then restarts, causing increased time for the router to be booted when a .scr file is used to boot up the router. [Defect ID 91685]
- ES2 10G LM fails the diagnostic test at "TenGe1Aps IOA Loopback Primary Phy PMA Test" when the Sumitomo 10KM-LR XFP transceiver is installed on the router. [Defect ID 91320]
- Unable to copy standby-disk0:system.log get error: Copy source does not exist or is unreachable. [Defect ID 90230]
- Outages occur with the SFM module. Fabric parity errors or fabric-to-slot link errors are not observed and all these counters show zero values. [Defect ID 90335]
- An error message is not displayed while the **copy running-configuration** command is used to copy the current configuration on the router to a script (*.scr) file. [Defect ID 90443]
- The state of the real-time clock battery cannot be tested and its working condition cannot be viewed using a **show** command. [Defect ID 90927]

TCP

- The security system logging message is generated at the info severity level with the label "couldn't allocate vty (no vtys are available)" even when virtual TTY (VTY) lines are available to be assigned for Telnet, SSH, or FTP services. [Defect ID 190135]

Unified ISSU

- When you perform a unified ISSU operation after link flaps in a network on a router that contains ES2 4G LMs and is configured with MPLS applications (L3VPNs, L2VPNs, and VPLS), all the ES2 4G LMs reset after the upgrade is completed. [Defect ID 189178]
- During the unified ISSU operation, if the router configuration is modified after the initialization phase of the process is completed and before the upgrade phase commences, the standby SRP module reloads for the second time when the synchronization process is in progress. After the second reload of the standby SRP module ends, the synchronization process also ends properly. [Defect ID 185517]

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 *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 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
 - Release 11.3.0

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
- The Routing Protocol Maximums table (for ERX310, ERX7xx, and ERX14xx routers) in *Appendix A, System Maximums*, of the *JunosE Release Notes*, for the following releases incorrectly mentions the maximum number of IP network interfaces (for IPv4 and IPv6) supported per chassis for the ERX310 router as 16,000. The correct maximum number of IP network interfaces supported per chassis for the ERX310 router is 32,000.
 - Releases 9.x, 10.x
 - Releases 11.0.0, 11.0.1, 11.0.2

- Releases 11.1.0, 11.1.1, 11.1.2
- Releases 11.2.0, 11.2.1
- Release 11.3.0

The table also incorrectly mentions the maximum number of IP network interfaces (for IPv4 and IPv6) supported per line module for all ERX routers as 8000. The correct maximum number of IP network interfaces supported per line module for all ERX routers is 16,383.

The correct values have been updated in the Routing Protocol Maximums table (for ERX310, ERX7xx, and ERX14xx routers) in *Appendix A, System Maximums*, of these *Release Notes*.

- The Routing Protocol Maximums table (for E120 and E320 routers) in *Appendix A, System Maximums*, of the *JunosE Release Notes*, for the following releases incorrectly mentions the maximum number of IP network interfaces (for IPv4 and IPv6) supported per line module for E120 and E320 routers.
 - Releases 9.x, 10.x
 - Releases 11.0.0, 11.0.1, 11.0.2
 - Releases 11.1.0, 11.1.1, 11.1.2
 - Releases 11.2.0, 11.2.1
 - Release 11.3.0

The following table lists the correct maximum number of IP network interfaces supported for the different line modules for E120 and E320 routers.

Routing Protocol Maximums Table (for E120 and E320 Routers)

Feature	E120	E320
IP network interfaces (IPv4 and IPv6)		
Per chassis	96,001	96,001
Per ES2 4G LM	16,383	16,383
Per ES2 10G LM	16,383	16,383
Per ES2 10G ADV LM	32,767	32,767
Per ES2 10G Uplink LM	16,383	16,383

The correct values have been updated in the Routing Protocol Maximums table (for E120 and E320 routers) in *Appendix A, System Maximums*, of these *Release Notes*.

- The Routing Protocol Maximums table (for E120 and E320 routers) in *Appendix A, System Maximums*, of the *JunosE Release Notes*, for the following releases incorrectly mentions the maximum number of MPLS next hops (egress FECs) when graceful restart is not enabled as 500,000 for all line modules on E120 and E320 routers:
 - Releases 9.x, 10.x
 - Releases 11.0.0, 11.0.1, 11.0.2
 - Releases 11.1.0, 11.1.1, 11.1.2
 - Releases 11.2.0, 11.2.1, 11.2.2
 - Releases 11.3.0, 11.3.1

The following table lists the maximum number of MPLS next hops (egress FECs) when graceful restart is not enabled for all line modules on E120 and E320 routers:

Routing Protocol Maximums table (for E120 and E320 routers)

Feature	E120	E320
IPv4 routing protocol scaling and peering densities		
MPLS next hops (egress FECs) when graceful restart is not enabled for ES2 4G LM	500,000	500,000
MPLS next hops (egress FECs) when graceful restart is not enabled for all line modules other than ES2 4G LM	300,000	300,000

- The Policy and QoS Maximums table (for ERX310, ERX7xx, and ERX14xx routers) in *Appendix A, System Maximums*, of the *JunosE Release Notes*, for the following releases fails to mention the maximum number of policy rules supported by the SRC client running on ERX routers:
 - Releases 11.2.0, 11.2.1, 11.2.2
 - Releases 11.3.0, 11.3.1, 11.3.2
 - Releases 12.0.0, 12.0.1, 12.0.2

Policy and QoS Maximums table (for ERX310, ERX7xx, and ERX14xx routers)

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Policy rules supported by the SRC client	256,000	256,000	256,000	256,000

Also, the following additional information applies to the allocation of memory blocks for policy rules sent to the SRC client from the SRC server for enforcing policy decisions:

For each rule that is sent from the SRC server using COPS messages to the SRC client, which is a router running JunosE Software, an entry is created in the policy table of the SRC client. A portion of the memory on the SRC client is needed to hold these policy rule entries that are transmitted to the SRC client for enforcing the policy decisions that are sent from the SRC server. The maximum number of memory blocks that is allocated to the SRC client functioning on the router for the policy rules that are sent from the SRC server is 256,000.

- The Policy and QoS Maximums table (for E120 and E320 routers) in *Appendix A, System Maximums*, of the *JunosE Release Notes*, for the following releases, fails to mention the maximum number of policy rules supported by the SRC client running on E120 and E320 routers:
 - Releases 11.2.0, 11.2.1, 11.2.2
 - Releases 11.3.0, 11.3.1, 11.3.2

The following table lists the maximum number of policy rules supported by the SRC client on E120 and E320 routers:

Policy and QoS Maximums table (for E120 and E320 routers)

Feature	E120	E320
Policy rules supported by the SRC client	256,000	256,000

Also, the following additional information applies to the allocation of memory blocks for policy rules sent to the SRC client from the SRC server for enforcing policy decisions:

For each rule that is sent from the SRC server using COPS messages to the SRC client, which is a router running JunosE Software, an entry is created in the policy table of the SRC client. A portion of the memory on the SRC client is needed to hold these policy rule entries that are transmitted to the SRC client for enforcing the policy decisions that are sent from the SRC server. The maximum number of memory blocks that is allocated to the SRC client functioning on the router for the policy rules that are sent from the SRC server is 256,000.

- The following JunosE documentation set fails to mention additional information regarding the creation of core dump files when the NMI button is depressed:
 - *ERX Hardware Guide—Resetting Line Modules and SRP Modules* section in *Chapter 9, Troubleshooting*
 - *E120 and E320 Hardware Guide—Resetting Line Modules and SRP Modules* section in *Chapter 9, Troubleshooting*

If you configure the router to generate core dumps using the **exception dump** command in Global Configuration mode, and if sufficient space is available at the destination to which the core dump files are transferred (such as the local flash card if you used the **local** keyword with the **exception dump** command), depressing the NMI button causes a core dump file to be stored on the flash card. This method of generating core dump files is applicable for both SRP modules and line modules. You can use the **show exception dump** command to determine whether core dumps are generated for only SRP modules or line modules.

- 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 Capability row in the Service Module (SM) Module Combination table of the *ERX Module Guide* fails to state the following additional information regarding the throughput for tunnel-server traffic:

Service Modules on ERX7xx models, ERX14xx models, and the ERX310 router can handle up to 800 Mbps of traffic, depending on the size of the packets. The throughput might be less with packets of smaller sizes.

- The Service Modules table in *ERX Module Guide, Appendix A, Module Protocol Support* incorrectly states that IPv6 Neighbor Discovery is not supported on the Service line module. In fact, IPv6 Neighbor Discovery is supported on ERX routers that contain a Service line module.
 - 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 Description row in the LM-10 ADV Line Module table of the *E120 and E320 Module Guide* fails to state that the ES2 10G ADV line module can also be paired with the ES2-S1 Service IOA.
 - 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.
 - The Capability row in the LM-4 Line Module table of the *E120 and E320 Module Guide* fails to state the following additional information regarding the throughput for tunnel-server traffic:

An ES2 4G LM with an ES2-S1 Service IOA (dedicated tunnel-server port) can receive traffic ranging from 3.5 Gbps (with 256 byte packets) to 3.8 Mbps (1024 to 1492 byte packets) for bidirectional L2TP LNS throughput. The throughput might be less with packets of smaller sizes.

A shared tunnel-server port on an ES2 4G LM can handle a maximum of 0.8 to 0.9 Gbps throughput for tunnel services, depending on the packet sizes. The throughput might be less with packets of smaller sizes.

- The Ethernet IOAs with ES2 10G ADV LM table in *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* erroneously indicates that tunnel-server ports are not supported for any Ethernet IOA combination with the ES2 10G ADV LM. Tunnel-server ports for shared tunnel services are supported on all Ethernet IOA combinations with the ES2 10G ADV LM.
- In the Module Combinations table in *E120 and E320 Module Guide, Appendix B, Module and Slot Combinations*, the ES2-S1 SERVICE row incorrectly states that the ES2-S1 Service IOA module cannot be paired with the ES2 10G ADV LM. The ES2-S1 Service IOA module is compatible with the ES2 10G ADV LM, in addition to the ES2 4G LM, for dedicated tunnel-server port functionality.
- The *Supported Modules for Dedicated Tunnel-Server Ports* section in *JunosE Physical Layer Configuration Guide, Chapter 6, Managing Tunnel-Service and IPSec-Service Interfaces* fails to mention that E120 and E320 routers support the ES2-S1 Service IOA and ES2 10G ADV LM combination and that you can install the ES2-S1 Service IOA with the ES2 10G ADV LM to use dedicated tunnel-server ports.
- The *ES2 10G Uplink LM Combination, ES2 10G LM Combination, and ES2 10G ADV LM Combination* sections in *JunosE Physical Layer Configuration Guide, Chapter 5, Configuring Ethernet Interfaces* incorrectly state that the maximum size of jumbo Ethernet frames supported at the ingress and egress interfaces is 9600 bytes. However, this maximum size applies only to frames that are received on the ingress interface or maximum receive unit (MRU). This maximum size does not apply to frames transmitted from the egress interface or maximum transmission unit (MTU). The following statement correctly describes the maximum size of jumbo Ethernet frames supported at the ingress and egress interfaces:
The IOA can also support 9600-byte jumbo Ethernet frames at the ingress (MRU) and 9188-byte jumbo Ethernet frames at the egress (MTU).
- The note in the *Configuring Tunnel-Server Ports and Tunnel-Service Interfaces* section in the *JunosE Physical Layer Configuration Guide, Chapter 6, Managing Tunnel-Service and IPSec-Service Interfaces* and the *Configuring QoS for Tunnel-Server Ports for L2TP LNS Sessions* section in the *JunosE Quality of Service Configuration Guide, Chapter 22, Configuring QoS for L2TP Sessions* does not accurately describe QoS profile support on interfaces stacked on the server-port. The following note accurately describes the QoS profile support:



Informational Note: Dedicated and shared tunnel-server ports on the ES2 10G ADV LM do not support QoS profiles on server-port interfaces and IP floating interfaces (IP interfaces that stack over MPLS stacked tunnels). However, you can configure QoS profiles for dedicated and shared tunnel-server ports on ES2 10G ADV LMs on interfaces other than server-port interfaces (such as ATM or Ethernet). On ES2 10G ADV LMs, you can also configure QoS profiles for dedicated and shared tunnel-server ports for L2TP LNS sessions only on interface types other than the server-port interface.

- The note in the *SMs, ES2-S1 Service IOA, and Shared Tunnel-Server Modules* section in *JunosE Physical Layer Configuration Guide, Chapter 6, Managing Tunnel-Service and IPSec-Service Interfaces* fails to state that in addition to ERX7xx models, ERX14xx models, and the ERX310 router, E120 and E320 routers also perform load balancing across all available server ports of the same type when both dedicated tunnel-server ports (on SMs) and shared tunnel-server ports (on shared tunnel-server modules) are configured.
- 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.
- In the *Relationship Between address and network Commands* section in *JunosE IP, IPv6, and IGP Configuration Guide, Chapter 4, Configuring RIP*, the **network** command section incorrectly states that you can specify either the standard subnet mask or the inverse subnet mask with the **network** command. The correct behavior of the **network** command is that you can specify only the standard subnet mask with this command. An error message is displayed if you attempt to configure an inverse subnet mask with this command.

Also, in this section, Example 2 describes the following sample usage of the **network** command with an inverse subnet mask:

```
host1(config-router)#network 10.2.1.0 0.0.0.255
jumbo
```

This example is incorrect. You cannot configure an inverse subnet mask with the **network** command.

- 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.
- *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.
- 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 output of the **show tunnel-server config** command in the *Identifying the Physical Location of the Tunnel-Server Port* section in *Chapter 6, Managing Tunnel-Service and IPSec-Service Interfaces*, in the *JunosE Physical Layer Configuration Guide* incorrectly displays the adapter identifier for the shared tunnel-server port configured on slot 12 and port 0.

In the last line of the output example, which displays Port 12/0/0 under the Port column, the adapter identifier is erroneously specified as 0. Because shared tunnel-server ports reside on a virtual adapter that is identified in the software as adapter 2, the correct slot/adapter/port combination for this configured shared tunnel-server is 12/2/0.

The correct output of the **show tunnel-server config** command is as follows:

```
host1#show tunnel-server config

                                Server Ports
                                -----
Port      Type      MaximumInterfaces  Provisioned  HwPresent  Bandwidth-
-----  -
Port 2/2/0 shared      8000           0           yes        N/A
Port 8/0/0 dedicated  16000          8000        yes        N/A
Port 12/2/0 shared      8000           0           yes        90
```

- 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 following additional information about the **pppoe motm** command is also inadvertently omitted:

- Use to cause the PPPoE application to send a PADM Message Of The Minute (MOTM) message. The recipient of the message is determined by the mode from which the command is issued.
- Use in Privileged Exec mode, Interface Configuration mode, and Profile Configuration mode.
- 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.

- The *Monitoring Redundancy in Installed Hardware* section in *JunosE Service Availability Configuration Guide, Chapter 2, Managing Module Redundancy* fails to state the following additional information regarding the **show environment all** command that is used to display system environment information for an E320 router.

You can use the **show environment all** command for ES2 4G line modules configured with Service IOAs on E120 and E320 routers to view information related to line module high availability configuration.

The following example displays information about the slots in which primary and secondary line modules that are configured and activated for stateful line module switchover reside. The information that relates to stateful line module switchover is highlighted in bold and appears in the line after the line module redundancy detail.

```
host1#show environment all
chassis: 17 slot (id 0x3, rev. 0x0)
fabric: 100 Gbps (rev. 1)
fans: fanSubsystemOk
nvs: ok (977MB flash disk, 29% full), matches running config
power: A ok, B not present
srp redundancy: mode is file-system-synchronization      auto-sync
enabled, switch-on-error enabled
status unknown
*** slots: cards missing or offline
online: 0 6 13
```

```

offline: 7
empty: 1 2 3 4 5 11 12 14 15 16
fabric slots: ok
online: 6 7 8 9 10
line redundancy: none
line card HA:
  high availability is configured on slots:
    {0, 13}
  high availability state is active on slots:
    {0, 13}
temperature: ok
timing: primary
primary: internal SC oscillator (ok)
secondary: internal SC oscillator (ok)
tertiary: internal SC oscillator (ok)
auto-upgrade enabled
fabric redundancy: ok

*** system operational: no

```

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

```

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

```

The descriptions of the fields related to line module high availability are as follows:

- line card HA—Configuration details for stateful line module switchover on E120 and E320 routers:
 - high availability is configured on slots—Slot numbers of the pairs of primary and secondary line modules configured on a single chassis for stateful switchover.
 - high availability state is active on slots—Slot numbers of the pairs of primary and secondary line modules activated for stateful switchover.
- The following additional information regarding the processing of tunneled subscriber sessions on Agent Circuit Identifier-based (ACI) VLAN subinterfaces with ICR partitions configured applies to the following guides in the JunosE documentation set:
 - *JunosE Broadband Access Configuration Guide—Monitoring Subscriber Information* section in *Chapter 2, Monitoring and Troubleshooting Remote Access*
 - *JunosE Service Availability Configuration Guide—Grouping ICR Subscribers Based on S-VLAN IDs* section in *Chapter 7, Managing Interchassis Redundancy*

If you attempt to bring up tunneled subscribers on ACI-based VLAN subinterfaces on LAC devices with subscriber groups that are based on S-VLAN IDs (using the **ip vrrp vrid icr-partition group svlan** command on S-VLAN subinterfaces), the VLAN subinterface does not come up and a log message to denote its down state is not generated. If you attempt to bring up tunneled subscribers on ACI-based VLAN subinterfaces on LAC devices with subscriber groups that are based on VLAN IDs (using the **ip vrrp vrid icr-partition group vlan** command on VLAN subinterfaces), the subscribers over tunnels are brought up. However, on the LAC device, the subscribers are logged in outside of the ICR partition.

This behavior is expected when attempts are made to log in tunneled subscribers over ACI-based VLAN subinterfaces configured with ICR partitions with VLAN-based grouping or S-VLAN based grouping.

- The *Specifying an OSPF Router ID* section in *Chapter 5, Configuring OSPF* in the *JunosE IP, IPv6, and IGP Configuration Guide* incorrectly specifies that the configuration mode for the **router-id** command in the example is Interface Configuration mode. This command is available only in Router Configuration mode. The correct example for this command is as follows:

```
host1(config-router)#router-id 192.168.50.5
```

Also, the **router-id** command section fails to state the following additional note about the function of this command for OSPFv2:

Although you can specify the router IP address using the **ip router-id** command in Global Configuration mode for OSPFv2 interfaces, use the **router-id** command in Router Configuration mode to enable the router to use a different IP address as the OSPF router ID rather than the address used for other IP routing protocols.

- The *Compressed IPv6 Formats* table in the *IPv6 Overview* section in *JunosE IP, IPv6, and IGP Configuration Guide, Chapter 2, Configuring IPv6* incorrectly states the IPv6 multicast address type. The correct IPv6 multicast address type is as follows:
 - Full Format—FF01:0:0:0:0:0:1F
 - Compressed Format—FF01::1F

- The *Configuring an IPv6 License* section in *JunosE IP, IPv6, and IGP Configuration Guide, Chapter 2, Configuring IPv6* erroneously specifies that you must configure an IPv6 license before you can use any IPv6 commands on the E Series router. The following information correctly describes IPv6 license configuration:

You must configure an IPv6 license before you can use any IPv6 routing protocol-related configuration commands on the E Series router. You can configure IPv6 policy settings, such as policy lists, classifier lists, rate-limit profiles, and local address pools for IPv6 subscribers even without configuring an IPv6 license.

- The *Removing an IPv6 Configuration* section in *JunosE IP, IPv6, and IGP Configuration Guide, Chapter 2, Configuring IPv6* incorrectly states that you can issue the **no ipv6** command to remove an IPv6 configuration from the virtual router. The following note describes the correct behavior of the **no ipv6** command:

You can remove all available IPv6 routing protocol-related configurations from the virtual router by using the **no ipv6** command. The previously configured IPv6 policy settings, such as policy lists, classifier lists, rate-limit profiles, and local address pools for IPv6 subscribers are not removed because these parameters can be configured without configuring an IPv6 license.

- In the *Using Expressions for Bandwidth and Burst Values in a Scheduler Profile* section in *JunosE Quality of Service Configuration Guide, Chapter 6, QoS Scheduler Hierarchy Overview*, the formulas used to translate burst values from bytes to milliseconds (ms), and to calculate the shaping rate and burst rate are incorrect. The following information describes the correct computation process for time in ms, shaping rate, and burst rate:

When calculating constant shaping rates, use the following formula to translate burst values from bytes to milliseconds (ms):

$$\text{Time (ms)} = [(\text{burstValueBytes} * 8 \text{ bits/byte}) / \text{Rate (bps)}] * 1000 \text{ (ms/s)}$$

Using this formula, a 2 Mbps service with a 500 KB burst yields:

$$[(500000 * 8) / 2000000] * 1000 = 2000 \text{ ms}$$

The shaping rate is calculated when the QoS profile is attached based on the parameter instance. For example:

```
host1(config)# scheduler-profile sp-1mbs
(config-scheduler-profile)# shaping-rate video-bandwidth % 100 burst 2000 milliseconds
```

When the shaping rate for video-bandwidth is 2 Mbps, the burst value is calculated using the following formula:

$$\text{Burst Value (bits)} = \text{Rate (bps)} * [\text{Time (ms)} / 1000 \text{ (ms/s)}]$$

The burst value in bits is calculated as:

$$\text{Burst Value (bits)} = 2000000 * [2000 / 1000] = 4000000$$

The burst value in bytes is calculated as:

$$\text{Burst Value (bytes)} = 4000000 / 8 = 500000$$

- 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 NBMA 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 *Duplicate and Broadcast Accounting* section in *JunosE Broadband Access Configuration Guide, Chapter 1, Configuring Remote Access* incorrectly states that the accounting information continues to be sent to the authenticating virtual router, but not to the operational virtual router, regardless of whether duplicate or broadcast AAA accounting is enabled on those routers. The following information describes the correct transmission of accounting information to authenticating virtual routers and operational virtual routers:

The accounting information is always sent to the authenticating virtual router. The accounting information is sent to the operational virtual router only if duplicate accounting is not enabled and if the authenticating virtual router is different than the operational virtual router.

- In the Juniper Networks (Vendor ID 4874) VSA Formats table in *JunosE Broadband Access Configuration Guide, Chapter 6, RADIUS Attribute Descriptions*, the description and value of the Med-Dev-Handle VSA attribute [26-59] are incorrect.

The correct format of the Med-Dev-Handle VSA is defined as follows:

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-59]	Med-Dev-Handle	Hexadecimal string used to determine mirror header attributes, prepended to each mirrored packet that is sent to the analyzer device	len	sublen	Salt encrypted string; hexadecimal string of 4 bytes or 8 bytes

- The Juniper Networks (Vendor ID 4874) VSA Formats table in *Chapter 6, RADIUS Attribute Descriptions* of the *JunosE Broadband Access Configuration Guide* inadvertently fails to mention the RADIUS VSAs numbered 26-101 through 26-105, and their associated definitions and formats.

See *JunosE 11.3.2 Release Notes, Appendix B, RADIUS VSA Formats* for detailed information

- The *L2TP Module Requirements* section in *JunosE Broadband Access Configuration Guide, Chapter 11, L2TP Overview* fails to state the following additional information regarding the available bandwidth for shared tunnel-server applications:

When you configure the GE-2 line module or the GE-HDE line module with a shared tunnel-server port, the available bandwidth for tunnel services is limited to 0.5 Gbps per module. When you configure the ES2 4G line module with a shared tunnel-server port, the available bandwidth for tunnel services is limited to 0.8 Gbps per module.

- The introductory paragraph in the *E120 Router and E320 Router* section of *Chapter 11, L2TP Overview, L2TP Module Requirements* of the *JunosE Broadband Access Configuration Guide* is ambiguous. The correct information is as follows:

To use an LNS on an E120 router or an E320 router, you must install an ES2 4G line module (LM) or an ES2 10G ADV LM with an ES2-S1 Service I/O adapter (IOA). With the ES2 4G LM, it is also possible to use an LNS with an IOA that supports the use of shared tunnel-server ports. For information about installing modules in these routers, see the *E120 and E320 Hardware Guide*.

- The *Monitoring DHCP External Server Statistics* section in *JunosE Broadband Access Configuration Guide, Chapter 22, Monitoring and Troubleshooting DHCP* incorrectly states under the Purpose section that you can use the **show ip dhcp-external statistics** command to display statistics for a specific server. Also, the command in the output example incorrectly states the following:

```
host1#show ip dhcp-external statistics server-address 100.20.32.1
```

The **server-address** keyword is not available for configuration with the **show ip dhcp-external statistics** command in Privileged Exec mode. Therefore, the output of this command displays statistics for all external servers configured on the router. The correct example for this command is as follows:

```
host1#show ip dhcp-external statistics
```

However, the sample output presented in the corresponding section of the guide is correct. When you issue the **show ip dhcp-external statistics** command, information about all configured DHCP external servers is displayed.

- 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.
- The Note in the **ipv6** command section in the *JunosE Command Reference Guide A to M* incorrectly states that disabling IPv6 on the router disables all IPv6 unicast routing protocols. The following note describes the correct processing of the **no ipv6** command:

Disabling IPv6 on the router causes only the IPv6 routing protocol-related configurations to be removed. The previously configured IPv6 policy settings, such as policy lists, classifier lists, rate-limit profiles, and local address pools for IPv6 subscribers are not removed because these parameters can be configured without configuring an IPv6 license.

- The syntax of the **mpls ldp ip-forwarding** command in the *JunosE Command Reference Guide A to M* incorrectly specifies that the **access-list**, **host-only**, and **prefix-list** keywords are optional arguments. The incorrect syntax of the command is as follows:

```
[ no ] mpls ldp ip-forwarding [ { access-list | prefix-list } listName ) ] [ host-only ]
```

The correct syntax of the command is as follows:

```
[ no ] mpls ldp ip-forwarding { access-list | prefix-list } listName host-only
```

Also, the Note in the command section incorrectly states that, although the carriage return, <cr>, option is displayed when you type a question mark (?) after entering the **mpls ldp ip-forwarding** command, this command takes effect only if the **access-list**, **host-only**, or **prefix-list** keyword, or a combination of them is specified. Additionally, it mentions that although the command is configured successfully and no error is displayed, labels are not used in the IP routing table for forwarding plain IP traffic.

The correct configuration behavior of this command is that you must enter one of the two keywords, specifically, **access-list** or **prefix-list**, along with the **mpls ldp ip-forwarding** command and the **host-only** keyword. Either of these keywords is required to enable labels to be used in the IP routing table for forwarding plain IP traffic. An error message is displayed when you press **Enter** to add LSPs to the routing table without specifying the access list or prefix list.

- The description of the **ppp initiate-ip** command in the *JunosE Command Reference N to Z Guide* fails to state the following information about this command:

Passive PPP clients are those subscribers configured with passive mode on dynamic or static PPP interfaces using the **ppp passive-mode** command. By default, passive mode is enabled on a PPP interface. Passive mode causes the PPP interfaces to wait for a period of one second before initiating LCP negotiation. This waiting period enables slow clients to start up and initiate LCP negotiation. Active PPP clients initiate the LCP negotiation process without waiting for the other side of the connection to initiate the negotiation.

If you configure a PPP interface without an IP interface or profile by not entering the **ppp initiate-ip** command, the router performs LCP negotiation for 2 to 3 minutes for passive clients. LCP negotiation is terminated after this period. To enable LCP negotiations to continue to occur for passive clients, you must enter the **ppp initiate-ip** command to enable PPP to create IP instances.

- The description of the **ppp initiate-ipv6** command in the *JunosE Command Reference N to Z Guide* fails to state the following information about this command:

Passive PPP clients are those subscribers configured with passive mode on dynamic or static PPP interfaces using the **ppp passive-mode** command. By default, passive mode is enabled on a PPP interface. Passive mode causes the PPP interfaces to wait for a period of one second before initiating LCP negotiation. This waiting period enables slow clients to start up and initiate LCP negotiation. Active PPP clients initiate the LCP negotiation process without waiting for the other side of the connection to initiate the negotiation.

If you configure a PPP interface without an IPv6 interface or profile by not entering the **ppp initiate-ipv6** command, the router performs LCP negotiation for 2 to 3 minutes for passive clients. LCP negotiation is terminated after this period. To enable LCP negotiations to continue to occur for passive clients, you must enter the **ppp initiate-ipv6** command to enable PPP to create IPv6 instances.

- The following note regarding the processing of the Ipv6-NdRa-Prefix RADIUS VSA attribute [26-129] applies to the description of the **ipv6-nd-ra-prefix** keyword available with the **radius include** command in the *JunosE Command Reference N to Z Guide*:

When you attempt to configure the Ipv6-NdRa-Prefix attribute using the dynamic configuration manager (DCM) profile, the prefix is not successfully configured and the subscriber does not come up. In this scenario, the RADIUS server rejects the authentication request from the subscriber and records an error message stating that address allocation failed. However, if you attempt to configure the Ipv6-NdRa-Prefix attribute using the RADIUS profile, the prefix is correctly configured and the subscriber comes up successfully. This behavior is expected when the DCM profile is used to configure the Ipv6-NdRa-Prefix attribute.

This scenario occurs when router advertisements are enabled in the DCM profile and the RADIUS server returns only the Framed-Interface-Id attribute. Because the AAA server requires one of the following attributes to authenticate IPv6 subscribers, and none of these attributes are returned from the RADIUS server, the logging in of subscribers fails:

- Ipv6-NdRa-Prefix (VSA 26-129)
- Framed-IPv6-Prefix (RADIUS IETF attribute 97)

- Framed-IPv6-Route (RADIUS IETF attribute 99)
- Framed-IPv6-Pool (RADIUS IETF attribute 100)
- Delegated-IPv6-Prefix (RADIUS IETF attribute 123)
- 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 Command Reference Guide N to Z*, the following additional information for the *queueDrainRate* variable in the **snmp-server host** command is missing:
By default, there is no limit on the number of traps sent per second to the host.

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	<i>General System Maximums</i> on page 70
Physical layer values	<i>Physical and Logical Density Maximums</i> on page 71
Link layer values	<i>Link Layer Maximums</i> on page 74
Routing protocol and performance values	<i>Routing Protocol Maximums</i> on page 79
Policy and QoS values	<i>Policy and QoS Maximums</i> on page 82
Tunneling values	<i>Tunneling Maximums</i> on page 85
Subscriber management values	<i>Subscriber Management Maximums</i> on page 87

System Maximums for E120 and E320 Routers	Section
General router values	<i>General System Maximums</i> on page 90
Physical layer values	<i>Physical and Logical Density Maximums</i> on page 91
Link layer values	<i>Link Layer Maximums</i> on page 93
Routing protocol and performance values	<i>Routing Protocol Maximums</i> on page 98
Policy and QoS values	<i>Policy and QoS Maximums</i> on page 101
Tunneling values	<i>Tunneling Maximums</i> on page 105
Subscriber management values	<i>Subscriber Management Maximums</i> on page 107

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	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 71.)				
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 71.)	4	–	–	14/24
Gigabit Ethernet ports per chassis (GE-8 I/O modules) (See Note 3 on page 71.)	4/16	–	–	14/96
Gigabit Ethernet ports per chassis (OC3-2 GE APS I/O module) (See Note 4 on page 71.)	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 EIs per chassis	504	1260	3024	3024
Logical E3s per chassis	24	60	144	144
Logical fractional EIs (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 EIs per cOCx/STMx F0 line module	252	252	252	252
	63 per OC3/STM1	63 per OC3/STM1	63 per OC3/STM1	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 74.)	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 74.)	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	32,768	81,920	96,000	96,000
(See Note 2 on page 74.)				
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				
	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

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
OCx/STMx POS	4	4	4	4
OC48 (ERX1440 router only)	–	–	–	1
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 <i>JunosE Link Layer Configuration Guide</i> .			
PPP interfaces per chassis (See Note 1 on page 74.)	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

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
PPPoE subinterfaces				
Subinterfaces per chassis (See Note 1 on page 74.)	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
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 per chassis (See Note 1 on page 79.)				
	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 79.)	32,000	32,000	32,000	48,000
Per line module	16,383	16,383	16,383	16,383
IPv4 routing protocol scaling and peering densities (See Note 3 on page 79.)				
Routing table entries (See Note 4 on page 79.)	500,000	500,000	500,000	500,000
ANCP Adjacency Scaling (See Note 5 on page 79.)	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)	1,000,000	1,000,000	1,000,000	1,000,000
MPLS next hops (egress FECs)	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
OSPF adjacencies	1000	1000	1000	1000
OSPF routes	25,000	25,000	25,000	25,000

Table 4: Routing Protocol Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
IPv6 routing table entries (See Note 3 on page 79.)	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 79.)	8000	8000	8000	8000
Total Martini circuits per chassis (See Note 7 on page 79.)	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 79.)	16,384	16,384	16,384	16,384
Outgoing interfaces per chassis (See Note 9 on page 79.)	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 80.)	400,000	400,000	400,000	400,000
Dynamic extended translations (NAPT) per SM (See Notes 10 and 11 on page 80.)	200,000	200,000	200,000	200,000
Response Time Reporter simultaneous operations per VR	500	500	500	500
VRRP VRIDs per line module	See <i>Ethernet VRRP VRIDs per line module</i> on page 76.			

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.
4. For each rule that is sent from the SRC server using COPS messages to the SRC client, which is a router running JunosE Software, an entry is created in the policy table of the SRC client. A portion of the memory on the SRC client is needed to hold these policy rule entries that are transmitted to the SRC client for enforcing the policy decisions that are sent from the SRC server. The maximum number of memory blocks that is allocated to the SRC client functioning on the router for the policy rules that are sent from the SRC server is 256,000.

Table 5: Policy and QoS Maximums

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
QoS queues per 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 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	256,000	256,000	256,000	256,000

(See Note 1 on page 82.)

Table 5: Policy and QoS Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Policy rules supported by the SRC client (See Note 4 on page 82.)	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 82.)	8127/16,255	8127/16,255	8127/16,255	8127/16,255
CAM entries (See Note 3 on page 82.)				
GE-2	64,000	–	–	64,000
GE-HDE	64,000	–	–	64,000
OC48/STM16	–	–	–	128,000
Policy egress interface attachments per 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 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 line module	24,575	24,575	24,575	24,575
Ingress per line module	24,575	24,575	24,575	24,575
Policy statistics blocks				
Egress per line module	256,000	256,000	256,000	256,000
Ingress per line module	256,000	256,000	256,000	256,000
Parent groups per line module				
GE-2, GE-HDE, and OC3/OC12 ATM line modules (Egress and Ingress)	24,575	24,575	24,575	24,575
All other line modules (Egress and Ingress)	8191	8191	8191	8191

Table 5: Policy and QoS Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Software lookup blocks				
Per line module	16,383	16,383	16,383	16,383
Secure policies (for packet mirroring)				
Per 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 85.)				
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 85.)				
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 85.)	16,000	16,000	16,000	32,000
L2TP sessions per line module (See Notes 1 and 3 on page 85.)				
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 85.)				
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 87.)	100,000	100,000	100,000	100,000
DHCP local server (See Note 2 on page 87.)				
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 87.)	32,000	32,000	32,000	32,000
DHCP relay and relay proxy client (See Notes 1 and 2 on page 87.)				
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 87.)				
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



Informational Note: The system maximum and line card maximum values mentioned in the tables are for single dimension scaling only. We recommend that you test scenarios which require scaling of multiple features to the maximum values concurrently, before deploying.

For example, on ERX1440 routers we support 48,000 PPP subscribers and 1,500,000 BGP 4 routes (NLRI). These values are independent of each other. We recommend that you test if the system can concurrently support 48,000 PPP subscribers and 1,500,000 BGP 4 routes (NLRI), before deploying.

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 90.)	3000	3000
Virtual routers and VRFs per line module, combined (See Notes 1 and 2 on page 90.)	3000	3000
ICR Partitions per chassis	640	640
CR 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 91.)		
10-Gigabit Ethernet ports per chassis (ES2-S1 10GE IOA)	6	12
10-Gigabit Ethernet ports per chassis (ES2-S2 10GE PR IOA)	6 + 6	12 + 12
(See Note 2 on page 91.)		
Gigabit Ethernet ports per chassis (ES2-S1 GE-4 IOAs)	24	48
Gigabit Ethernet ports per chassis (ES2-S1 GE-8 IOAs)	96	192
(See Note 2 on page 91.)		
Gigabit Ethernet ports per chassis (ES2-S3 GE-20 IOA)	120	240
(See Note 2 on page 91.)		
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 93.)	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 93.)	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 93.)	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 93.)	64,000	96,000
Bridged Ethernet interfaces per line module (OCx/STMx ATM)	16,000	16,000

Table 10: Link Layer Maximums Table continued

Feature	E120	E320
Dynamic interfaces		
Active autosensed dynamic interface columns per chassis over static or dynamic (bulk) ATM1483 subinterfaces (See Note 2 on page 93.)	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 93.)	64,000	96,000
Ethernet S-VLANs per IOA		
(See Note 6 on page 93.)		
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 93.)	64,000	96,000
Ethernet VLANs per IOA		
(See Note 7 on page 93.)		
ES2-S1 GE-4 IOA (with ES2 4G LM) (See Note 5 on page 93.)	16,384	16,384
ES2-S1 GE-8 IOA (with ES2 4G LM or ES2 10G LM) (See Note 5 on page 93.)	16,384	16,384
ES2-S1 GE-8 IOA (with ES2 10G ADV LM) (See Note 5 on page 93.)	32,768	32,768

Table 10: Link Layer Maximums Table continued

Feature	E120	E320
ES2-S1 10GE IOA (with ES2 4G LM) (See Note 5 on page 93.)	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 93.)	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 93.)		
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
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 <i>JunosE Link Layer Configuration Guide</i> .	

Table 10: Link Layer Maximums Table continued

Feature	E120	E320
PPP major interfaces per chassis (See Notes 2 and 3 on page 93.)	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 93.)	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 93.)	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
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. 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
IPv4 forwarding table entries per chassis (See Note 1 on page 98.)	1,048,576	1,048,576

Table 11: Routing Protocol Maximums Table continued

Feature	E120	E320
IP network interfaces (IPv4 and IPv6)		
Per chassis (See Note 2 on page 98.)	96,001	96,001
Per ES2 4G LM	16,383	16,383
Per ES2 10G LM	16,383	16,383
Per ES2 10G ADV LM	32,767	32,767
Per ES2 10G Uplink LM	16,383	16,383
IPv4 routing protocol scaling and peering densities		
(See Note 3 on page 98.)		
Routing table entries (See Note 4 on page 98.)	500,000	500,000
ANCP Adjacency Scaling (See Note 5 on page 98.)	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 for ES2 4G LM	500,000	500,000
MPLS next hops (egress FECs) when graceful restart is not enabled for all line modules other than ES2 4G LM	300,000	300,000
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	100,000	100,000
(See Note 3 on page 98.)		
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 98.)	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 98.)	16,384	16,384
Outgoing interfaces per chassis (See Note 8 on page 98.)	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 <i>Ethernet VRRP VRIDs per line module</i> on page 96.	See <i>Ethernet VRRP VRIDs per line module</i> on page 96.

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
 - 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.
 4. For each rule that is sent from the SRC server using COPS messages to the SRC client, which is a router running JunosE Software, an entry is created in the policy table of the SRC client. A portion of the memory on the SRC client is needed to hold these policy rule entries that are transmitted to the SRC client for enforcing the policy decisions that are sent from the SRC server. The maximum number of memory blocks that is allocated to the SRC client functioning on the router for the policy rules that are sent from the SRC server is 256,000.

Table 12: Policy and QoS Maximums

Feature	E120	E320
QoS queues per line module (See Note 1 on page 101.)	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

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
QoS scheduler nodes per line module	64,000	64,000
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 rules supported by the SRC client (See Note 4 on page 101.)	256,000	256,000
Policy egress interface attachments per line module (See Note 2 on page 101.)		
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 101.)		
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

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
ES2 10G LM IP 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	64,000	64,000
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 combined IP and IPv6 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

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
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
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 101.)	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 105.)	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 105.)	4000	4000
L2TP sessions per chassis (See Note 2 on page 105.)	60,000	60,000
L2TP sessions per line module with shared tunnel-server ports provisioned (See Note 2 on page 105.)	8000	8000
L2TP sessions per ES2-S1 Service IOA (See Note 2 on page 105.)	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 105.)	8000	8000
L2TP tunnels per ES2-S1 Service IOA (See Note 1 and Note 2 on page 105.)	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 107.)	100,000	100,000
DHCP local server (See Note 2 on page 107.)		
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 107.)	32,000	32,000
DHCP relay and relay proxy client (See Notes 1 and 2 on page 107.)		
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 107.)		
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



Informational Note: The system maximum and line card maximum values mentioned in the tables are for single dimension scaling only. We recommend that you test scenarios which require scaling of multiple features to the maximum values concurrently, before deploying.

For example, on E320 routers we support 96,000 PPP subscribers and 1,500,000 BGP 4 routes (NLRI). These values are independent of each other. We recommend that you test if the system can concurrently support 96,000 PPP subscribers and 1,500,000 BGP 4 routes (NLRI), before deploying.