

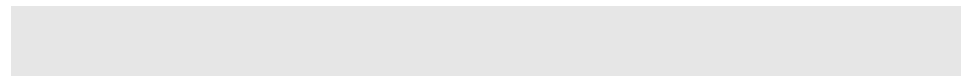


JunosE[™] Software for E Series[™] Broadband Services Routers

Release Notes

Release

14.2.0



Published: 2013-04-12

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 © 2013, Juniper Networks, Inc.

All rights reserved. Printed in USA.

JunosE™ Software for E Series™ Broadband Services Routers Release Notes, Release 14.2.0

Revision History

April 2013—FRS JunosE 14.2.0

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 14.2.0	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	2
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
AAA	7
L2TP	9
Policy Management	10
SDX Software and SRC Software	15
Service Manager	16
System Maximums	16
Unified ISSU	17
Unsupported Features	18
E120 Router and E320 Router	18
Policy Management	18
Stateful SRP Switchover (High Availability)	18
Release Software Protocols	19
Core Routing Stack	19
Layer 2 Protocols	19
Multiprotocol Label Switching (MPLS)	19
Network Management Protocols	19
Routing Protocols	19
Security Protocols	20
SRC Software and SDX Software Compatibility Matrix	20
Known Behavior	21
AAA	21
ATM	21
BGP	21
BGP/MPLS VPNs	22
B-RAS	22

CLI	22
DHCP.....	26
DHCP External Server	26
Dynamic Interfaces.....	27
Ethernet.....	27
Flash	28
GRE	28
Hardware.....	29
HDLC.....	29
IP.....	30
IPsec	31
IS-IS.....	32
L2TP	32
LDP	33
Line Module Redundancy.....	33
MLPPP	33
MPLS.....	33
Multicast	33
Packet Mirroring.....	35
Policy Management	35
PPP	38
PPPoE.....	38
QoS	38
RADIUS	38
SNMP	39
SRC Software and SDX Software	40
SSH	41
Stateful SRP Switchover (High Availability)	41
System.....	41
Tunneling	42
Known Problems and Limitations	43
ATM.....	43
BFD	43
DHCP.....	43
Forwarding.....	43
ICR	44
IGMP	44
IS-IS.....	45
L2TP	45
MLD.....	46
MPLS.....	46
Policy Management	46
QoS	47
SDX Software and SRC Software	47
Server Card Manager (SCM).....	48
Service Manager.....	48
Stateful Line Module Switchover (High Availability)	49
Stateful SRP Switchover (High Availability) and IP Tunnels	49
Subscriber Management	49
System.....	50

	TCP	50
	Resolved Known Problems	51
	Policy Management	51
	SNMP	51
	Errata	51
Appendix A	System Maximums	53
	ERX310, ERX7xx, and ERX14xx System Maximums	54
	General System Maximums	54
	Physical and Logical Density Maximums	55
	Link Layer Maximums	58
	Routing Protocol Maximums	63
	Policy and QoS Maximums	66
	Tunneling Maximums	69
	Subscriber Management Maximums	71
	E120 and E320 System Maximums	74
	General System Maximums	74
	Physical and Logical Density Maximums	75
	Link Layer Maximums	77
	Routing Protocol Maximums	82
	Policy and QoS Maximums	85
	Tunneling Maximums	89
	Subscriber Management Maximums	91

Release 14.2.0

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 about 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-numbered 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 14.2.0 of the system software for the Juniper Networks E Series Broadband Services Routers and contain the following sections:

- *Release Highlights* on page 6
- *Unsupported Features* on page 18
- *Release Software Protocols* on page 19
- *SRC Software and SDX Software Compatibility Matrix* on page 20
- *Known Behavior* on page 21
- *Known Problems and Limitations* on page 43
- *Resolved Known Problems* on page 51
- *Errata* on page 51
- *Appendix A, System Maximums*, on page 53

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 14.1.0 and 14.2.0. 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 14.2.0:

- These 14.2.0 *Release Notes*, which describe changes between Release 14.1.0 and Release 14.2.0
- The 14.1.0 *Release Notes*, which describe features available in Release 14.1.0
- The 14.2.x documentation set, which provides detailed information about features available in Release 14.2.0

The 14.2.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 website 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 14.2.0 includes the features described in this section.

Category	Feature
AAA	<ul style="list-style-type: none"> Support for Configuring Interim, Broadcast, and Policy-Based Accounting in PPP Profiles and Broadcast Virtual Router Groups on page 7
L2TP	<ul style="list-style-type: none"> Controlling the Transmission of L2TP Hello Messages to L2TP Peers on page 9
Policy Management	<ul style="list-style-type: none"> Support for Classifier-Specific Statistics Accounting Feature for IPv4 and IPv6 Interfaces on page 10 CLI Support for Monitoring Policy Resource Utilization and Generating SNMP Traps on Policy Resources Exhaustion on page 11 Support for Detection of Corruption in the Statistics FPGA for AAA-Based Policy Accounting on page 12 Support for Processing Source and Destination Route-Class Classifiers in Policies for Dynamic MPLS VPNs for ES2 10G, ES2 10G Uplink, and ES2 10G ADV LMs on page 14
SDX Software and SRC Software	<ul style="list-style-type: none"> Support for Selecting the Primary SRC Server After the Restart of an SRC Client for Applying Policies on page 15
Service Manager	<ul style="list-style-type: none"> Support for the Collection of IPv6 Service Accounting Statistics Separately Through IPv6 Accounting VSAs for Dual-Stack Subscribers on page 16

Category	Feature
System Maximums	<ul style="list-style-type: none"> Increased Number of Dual-Stack Subscribers over PPP on E120 and E320 Routers Without Chassis Cluster or Unified ISSU Support on page 16
Unified ISSU	<ul style="list-style-type: none"> Unified ISSU Support for OSPFv3 on page 17

AAA

- Support for Configuring Interim, Broadcast, and Policy-Based Accounting in PPP Profiles and Broadcast Virtual Router Groups

You can now use the **aaa virtual-router** command with the **interim-accounting** or **policy-accounting** keyword to enable or disable the option to send periodic interim accounting messages and policy-based accounting statistics from authentication, authorization, and accounting (AAA) to the broadcast accounting servers that are configured in a broadcast virtual router group.

When the sending of policy-based accounting statistics is disabled, AAA sends the user accounting statistics that are received from Point-to-Point Protocol (PPP) to the broadcast accounting servers. Policy-based accounting messages that are sent from AAA contains the Calling-Station-Id [31] and Event-Timestamp [55] RADIUS attributes.

The output of the **show aaa accounting vr-group** command contains the status of the interim-accounting and policy-accounting features for virtual routers that are added to the virtual router group.

You can now use the **aaa accounting interim-update** command to enable or disable the option to send periodic interim accounting messages at the configured user accounting interval to a primary accounting server on a per-virtual router basis. When this option is disabled, interim accounting messages are not sent to the primary accounting server even though the user accounting interval is configured. When the user accounting interval is set as 0, interim accounting messages are not sent to the primary accounting server even though the option to send periodic interim accounting messages is enabled.

The output of the **show aaa accounting** command contains the status of the option to periodically send interim accounting messages to the primary accounting server.

You can now use the **ppp aaa-accounting-broadcast** command to assign a broadcast virtual router group to a PPP profile, which enables broadcast accounting in the PPP profile. If the virtual router group exists in the router, AAA reads the virtual router group configuration and sends the broadcast accounting messages on the basis of the virtual router group configuration to the broadcast accounting servers.

When the broadcast virtual router group is configured at both the PPP profile and virtual router levels, AAA sends accounting messages only to broadcast accounting servers in the broadcast virtual router group that is configured at both the PPP profile and virtual router levels.

The output of the **show profile name profileName** command contains the name of the broadcast virtual router group that is assigned to the PPP profile.

The configurations done using the **aaa virtual-router**, **aaa accounting interim-update**, and **ppp aaa-accounting-broadcast** commands are saved across the chassis cluster and unified in-service software upgrade (ISSU). The following configurations are the default configurations:

- The option to send periodic interim accounting messages to the broadcast accounting servers is enabled. You can use the **no** version of the **aaa virtual-router** command with the **interim-update** keyword to restore this configuration.
- The option to send policy-based accounting messages to the broadcast accounting servers is disabled. You can use the **no** version of the **aaa virtual-router** command with the **policy-accounting** keyword to restore this configuration.
- The option to send periodic interim accounting messages to the primary accounting server is enabled. You can use the **no** version of the **aaa accounting interim-update** command to restore this configuration.



Informational Note: Policy-based accounting is not supported on the primary accounting server. IP version 6 (IPv6) policy-based accounting is not supported for the IPv6 policy returned from the RADIUS server.

The following commands have been added to support this feature:

- **aaa accounting interim-update**
- **ppp aaa-accounting-broadcast**

The following command has been enhanced to support this feature:

- **aaa virtual-router**

The output of the following commands have been modified to support this feature:

- **show aaa accounting**
- **show aaa accounting vr-group**
- **show profile**

As part of this feature, the `aaaServerGeneral` event log has been modified to log the following debug messages:

- Do not send interim accounting to the primary accounting server.
- Interim accounting is disabled on the broadcast virtual router.
- User accounting is disabled on the broadcast virtual router.

Change in existing behavior: New feature added as described here.

L2TP

- Controlling the Transmission of L2TP Hello Messages to L2TP Peers

In some cases, the Layer 2 Tunneling Protocol (L2TP) peer of an E Series router acting as an L2TP access concentrator (LAC) or L2TP network server (LNS) is incapable of responding with a Zero Length Body (ZLB) message to acknowledge the receipt of Hello messages sent by the LAC or LNS. This causes the LAC or LNS to terminate the control channel. You can now use the **l2tp disable tunnel-hello** command to disable the transmission of Hello messages from the LAC or LNS. This feature can be used in scenarios where there is no support available for the transmission of Hello messages so that unintentional control channel termination is avoided.



Informational Note: When the transmission of Hello messages is disabled, the following effects are expected:

If the PPP keepalive message is also disabled and the L2TP peer goes out of reach, the control connection cannot be closed at the LAC or LNS end because no explicit control message is received by the LAC or LNS from its peer. This causes the connection to use the LAC or LNS resources unnecessarily. If the maximum number of connections were in use before the peer went out of reach, the peer might not be able to reconnect after it comes up because all the L2TP resources are used up by the previous connections. The resultant behavior depends on the peer L2TP implementation.

JunosE Silent Failover (SF)–based control connection recovery is dependent on the transmission of Hello messages to synchronize sequence numbers with the L2TP peer after failover. With the transmission of Hello messages disabled, SF-based recovery cannot be performed.

You can now use the **no l2tp disable tunnel-hello** command to enable the transmission of Hello messages from the LAC or LNS.



Informational Note: Execution of the **no l2tp disable tunnel-hello** command resumes the transmission of Hello messages in the tunnel and results in the termination of the control connection if the L2TP peer did not support Hello messages.

The following command has been added to support this feature:

- **l2tp disable tunnel-hello**

Change in existing behavior: New feature added as described here.

Policy Management

- Support for Classifier-Specific Statistics Accounting Feature for IPv4 and IPv6 Interfaces

The JunosE Software now enables you to take into account statistics about IP version 4 (IPv4) and IP version 6 (IPv6) interfaces only for classifier groups that have the classifier action set as **aaa count enable**. You can enable the classifier-specific statistics accounting feature for a classifier group by using the **aaa count enable** command in Classifier Group Configuration mode.



Informational Note: The classifier-specific statistics accounting feature is not supported for secondary input policy and hierarchical rate-limit policy.

When the classifier-specific statistics accounting feature is enabled, only the IP or IPv6 and user payload headers are included in policy octet counters of upstream and downstream packet statistics. The control overheads are excluded from policy octet counters.

When the classifier-specific statistics accounting feature is enabled on any of the classifier groups assigned to IPv4 and IPv6 interfaces, the router adds the forwarded bytes of only the classifier groups that have the classifier action set as **aaa count enable** and sends the total bytes to authentication, authorization, and accounting (AAA) for service accounting.



Informational Note: To exclude Dynamic Host Configuration Protocol version 6 (DHCPv6) or multicast packets from service accounting, you must explicitly configure a corresponding classifier group that matches the DHCPv6 or multicast packets and for which the classifier-specific statistics accounting feature is disabled.

You can disable the classifier-specific statistics accounting feature for a classifier group by using the **no aaa count enable** command. By default, the classifier-specific statistics accounting feature is disabled for the classifier group.

The outputs of the **show ip interface** and **show ipv6 interface** commands are updated to include the **aaa count enable** field for the classifier group when you enable the classifier-specific statistics accounting feature.

The following command has been added to support this feature:

- **aaa count enable**

The output of the following commands have been enhanced to support this feature:

- **show ip interface**
- **show ipv6 interface**

Change in existing behavior: New feature added as described here.

- CLI Support for Monitoring Policy Resource Utilization and Generating SNMP Traps on Policy Resources Exhaustion

The maximum number of policies that can be attached to an interface on an E Series router depends on the classifier entries and the number of attachment resources available on the interface. JunosE Software allocates interface-attachment resources when policies are attached to the interfaces. You can now use the **show policy-resources slot** command to monitor policy resource consumption on a specific slot or on all slots to which the policies are attached. The **show policy-resources slot** command displays the following details:

- Type of the policy resource
- Number of policy resources supported by the line module for each resource type
- Number of policy resources consumed by the line module for each resource type
- Number of policy resources that are free in the line module for each resource type
- Direction type for each resource type



Informational Note: The current system maximums of the E Series router are displayed in the **show policy-resources slot** command output as the number of policy resources supported by the line module. For current system maximums, see [Appendix A, System Maximums in this release notes](#).

You can now use the **policy-resource-exhaustion trap enable** command to enable the policy resource exhaustion trap to send an SNMP trap notification when the policy resources are exhausted. The notification contains the following information:

- Slot number
- Resource type
- Resource direction type

You can view the status of the policy exhaustion trap by using the **show policy-resources trap** command.

As part of this feature, the following SNMP MIB objects are added to the Juniper Networks Policy MIB:

- rsPolicyResourceSlotNumber—Displays the slot number on which the policy resources are exhausted. The value can range from 0 through 255.
- rsPolicyResourcePoolType—Displays the policy resource pool type of the exhausted policy resources. This can be:
 - > ipPolicyDescriptor (0)
 - > l2PolicyDescriptor (1)
 - > rateLimiter (2)
 - > statsBlock (3)
 - > camTotalEntries (4)
 - > softwareLookupRuleSet (5)
 - > parentGroup (6)

- **rsPolicyResourceDirection**—Displays the policy resource pool direction of the exhausted policy resources. This can be:
 - > ingress (0)
 - > egress (1)
 - > both (2)

The following commands have been added to support this feature:

- **show policy-resources**
- **policy-resource-exhaustion trap enable**
- **show policy-resources trap**

Change in existing behavior: New feature added as described here.

- Support for Detection of Corruption in the Statistics FPGA for AAA-Based Policy Accounting

When a bit flip occurs in the RAM of the statistics field-programmable gate array (FPGA) of a router, the router may transmit incorrect accounting details for IP version 4 (IPv4) and IP version 6 (IPv6) subscribers to a RADIUS server through authentication, authorization, and accounting (AAA). You can solve the problem of inaccurate policy accounting details by replacing the defective hardware.

To avoid accounting wrong statistics before replacing the defective hardware, you can configure a software-detection mechanism. This mechanism compares the interface counters and policy counters against a configured threshold value to detect the statistics FPGA corruption before the router sends the accounting details to the RADIUS server.

If the difference between the interface counters and policy counters for ingress or egress policies collected over two polling intervals is equal or greater than the configured threshold value, a corruption is detected in the statistics FPGA. If the difference between the interface counters and policy counters for ingress or egress policies collected over two polling intervals is less than the configured threshold value, no corruption is detected in the statistics FPGA.

The software-detection mechanism already implemented to detect the statistics FPGA corruption for Session and Resource Control (SRC) server-based policy accounting is used to detect the statistics FPGA corruption for AAA-based policy accounting. You can use the following existing commands to configure the software-detection mechanism and view the configuration:

- **fpga-stats-monitoring-enable** - Enables the router to detect the corruption in the statistics FPGA.
- **fpga-stats-monitoring threshold** - Specifies the threshold value to be used to determine corruption in the statistics FPGA.
- **show fpga-stats-monitoring** - Displays the settings configured for the detection of the statistics FPGA corruption.

This detection mechanism is triggered on receiving the following requests from the RADIUS server:

- Accounting start request
- Interim accounting request

- Accounting stop request

This detection mechanism performs the corrections in the counters by using the following formulas while comparing the interface and policy counters:

- Policy byte counter = Sum of policy bytes – (Policy packet counter * Extra header)
- Egress policy packet counter = Policy packet counter – Out-policed packet counter
- Egress policy byte counter = Policy byte counter – Out-policed byte counter
- Policy byte counter (in multicast scenario) = Interface receive byte counter – Multicast byte counter
- Policy packet counter (in multicast scenario) = Interface receive packet counter – Multicast packet counter

When corruption is detected in the statistics FPGA, AAA logs out the affected subscribers, prevents new subscribers from establishing a session on the defective slot, and generates the system log message and SNMP trap containing the affected subscriber's interface name, username and defective slot information. The SNMP trap is sent to an external device. AAA maintains the defective slot information until the defective hardware is replaced.

If Line Card Redundancy (LCR) and Line Card High Availability (LCHA) are not enabled on the defective hardware, AAA does not disconnect the affected subscribers but prevents any new subscribers from being connected to the defective slot when the corruption is detected.

If LCR is enabled in the defective hardware, AAA terminates all affected subscribers and the standby slot takes over when the corruption is detected.

If LCHA is enabled in the defective hardware, LCHA group is unconfigured when the corruption is detected:

- If the configured primary Tunnel Service line module (TSM) is active, the standby TSM is reloaded to accept new subscribers while maintaining the subscribers existing in the primary TSM.
- If the configured secondary TSM is active, the primary and standby TSM cards are reloaded to terminate all subscribers.



Informational Note: The maximum number of supported interfaces should be set as zero on the current active TSM card to achieve load balancing.

If the accounting interim is enabled on the router, AAA would inform Point-to-Point Protocol (PPP) to poll accounting requests at the frequency of the configured accounting interim period when a subscriber logs in.

This feature has the following limitations:

- Information about defective slots is not persistent across switch route processor (SRP) resets. Therefore, if subscribers attempt to log in to a slot, which was determined to be corrupted before the reset, they are permitted to log in until the detection mechanism again classifies the slot to be defective after the reset.

- The detection mechanism has a limitation in the calculation of policy statistics when the ingress or egress traffic does not match any of the classified rules configured within a policy. To avoid this discrepancy, the unaccounted traffic should be matched with a default classifier group.
- The detection mechanism cannot detect bit flips in the least significant bits, which results in statistics corruption lower than the configured threshold value.
- The detection mechanism is not supported for the Hierarchical Rate Limit Profile (H-RLP).

The following command has been added to support this feature:

- **show aaa-corrupted-slots**

Change in existing behavior: Existing feature extended as described here. In lower-numbered releases, detection of corruption in the statistics FPGA is implemented for SRC server-based policy accounting. The new feature detects statistics FPGA corruption for AAA-based policy accounting and displays the corrupted slot details.

- Support for Processing Source and Destination Route-Class Classifiers in Policies for Dynamic MPLS VPNs for ES2 10G, ES2 10G Uplink, and ES2 10G ADV LMs

You can now configure source route-class and destination route-class classifier attributes in a classifier control list (CLACL) and attach a policy that references the CLACL to dynamic IPv4 or IPv6 interfaces over MPLS VPNs to enable the incoming and outgoing packets to be classified based on the route-class values. The route-class value that is obtained from the address lookup performed on the source address of a packet is referred to as the source route-class. The route-class value that is derived from the address lookup performed on the destination address of a packet is referred to as the destination route-class.

You can use the **ip classifier-list** or the **ipv6 classifier-list** commands in Global Configuration mode to configure the source address lookup route-class values and destination address lookup route-class values in a CLACL. When you create an IPv4 or an IPv6 policy list that contains a CLACL with a route-class value and the policy list is attached to a dynamic MPLS VPN interface, the incoming or outgoing packets that are enabled for the source or destination route-class are classified based on the route-class classifier and the action configured for the route-class classifier is taken for MPLS VPN packets.

This functionality of processing the route-class classifiers in CLACLs that are contained in input and output policy lists attached to dynamic MPLS VPN interfaces is supported on ES2 10G, ES2 10G Uplink, and ES2 10G ADV LMs.

Change in existing behavior: Existing feature extended as described here. In lower-numbered releases, you could configure source address and destination address lookup route-class classifier attributes in a CLACL to classify packets based on the configured route-class for only ES2 4G LMs and OCx/STMx ATM line modules with OC3-4 I/O modules. If you configured route-class classifiers in policies attached to dynamic IPv4 or IPv6 interfaces over MPLS VPNs for ES2 10G, ES2 10G Uplink, or ES2 10G ADV LMs, the incoming and outgoing packets were classified based on the action specified for the default classifier and the actions configured for the route-class classifier were not considered.

SDX Software and SRC Software

- Support for Selecting the Primary SRC Server After the Restart of an SRC Client for Applying Policies

When an E Series router acting as an SRC client is rebooted, it selects its SRC server randomly from the primary, secondary, or tertiary SRC servers without waiting to establish a connection with the primary SRC server. The SRC server that first establishes a connection with the SRC client is selected by the SRC client.

You can now use the newly introduced **sscc connectivityTimer** command to define a time period in seconds until which the rebooted SRC client waits for the primary SRC server to establish a connection with the SRC client. If the connection is not established within the configured time period, the SRC client selects the SRC server randomly, which is the default behavior of the SRC client. The connectivity timer value ranges from 0 through 300 seconds. By default, the timer value is 0 seconds, which denotes that the SRC client selects the SRC server randomly without waiting for the primary SRC server to establish a connection with the SRC client.

If the connectivity timer value is configured, the following functionalities are performed during the rebooting of the SRC client:

- The connectivity timer and ping operation to the primary SRC server are initiated when the first uplink or downlink line module other than the ES2 4G or ES2 10G ADV LMs with ES2-S1 Service IOA comes online.
- The ping operation to the primary SRC server is tried every 3 seconds after the timer is started until the timer stops or expires.
- If the ping operation is successful within the configured timer value, the timer is stopped and the SRC client selects the primary SRC server. If the ping operation is not successful until the expiry of the configured timer value, the ping operation is stopped and the SRC client selects the SRC server that first established a connection with the SRC client.

The following command has been added to support this feature:

- **sscc connectivityTimer**

The output of the following command has been enhanced to support this feature:

- **show sssc info**

As part of this feature, the ssscGeneral log has been modified to add the SDX connection replies and connectivity timer attributes at the debug level.

Change in existing behavior: Existing feature extended as described here. In lower-numbered releases, the SRC client selects the SRC server randomly after rebooting without waiting for the primary SRC server to establish a connection with the SRC client.

Service Manager

- Support for the Collection of IPv6 Service Accounting Statistics Separately Through IPv6 Accounting VSAs for Dual-Stack Subscribers

The JunosE Software now allows the router to send IPv6 service accounting statistics separately through IPv6 accounting vendor-specific attributes (VSAs) to the accounting server when the subscriber is either an IPv6 subscriber or a combination of IPv4 and IPv6 subscribers in a dual-stack environment. This feature is also applicable for a single-stack environment.

The Service Manager tracks IPv6 statistics by using the following JunosE objects that are passed back through the **env.setResult** method:

- secondary-input-stat-clacl
- input-stat-clacl
- output-stat-clacl
- input-stat-epg
- output-stat-epg

IPv6 service accounting statistics are sent to the accounting server through the Acct-Stop and Interim-Acct RADIUS messages by using the following VSAs:

- Ipv6-Acct-Input-Octets [26-151]
- Ipv6-Acct-Output-Octets [26-152]
- Ipv6-Acct-Input-Packets [26-153]
- Ipv6-Acct-Output-Packets [26-154]
- Ipv6-Acct-Input-Gigawords [26-155]
- Ipv6-Acct-Output-Gigawords [26-156]

The chassis cluster and unified in-service software upgrade (ISSU) are supported for this feature.



Informational Note: IPv6 service accounting statistics are sent through RADIUS messages only if the inclusion of IPv6 accounting RADIUS attributes in a RADIUS message is enabled by using the **radius include** command.

Change in existing behavior: Existing feature extended as described here. In lower-numbered releases, IPv6 user accounting statistics are sent separately through IPv6-specific VSAs, whereas IPv6 service accounting statistics are sent together with IPv4 service accounting statistics through IPv4-specific VSAs.

System Maximums

- Increased Number of Dual-Stack Subscribers over PPP on E120 and E320 Routers Without Chassis Cluster or Unified ISSU Support

The E120 and E320 router chassis-wide limit of dual-stack subscribers over Point-to-Point Protocol (PPP) is increased from 48,000 to 64,000 without chassis cluster or unified in-service software upgrade (ISSU) feature support. In addition, the following limits are increased for supporting 64,000 dual-stack PPP subscribers:

- Policy configuration per chassis is increased from 96,000 to 128,000.

- Quality-of-service (QoS) configuration per chassis is increased from 96,000 to 128,000.
- The total number of IP version 4 (IPv4) and IP version 6 (IPv6) dynamic interfaces per chassis is increased from 96,001 to 128,001.



Informational Note: Even though 128,001 IP interfaces are supported, only a maximum of 96,000 subscribers are supported per chassis. A combination of single-stack and dual-stack subscribers to use all 128,001 interfaces is not supported.

The chassis-wide limit of 64,000 dual-stack PPP subscribers applies only to the following interface stack:

- IPv4 over PPP / Point-to-Point Protocol over Ethernet (PPPoE) / Ethernet
- IPv6 over PPP / PPPoE / Ethernet
- IPv4 over PPP / PPPoE / virtual LAN (VLAN) / Ethernet
- IPv6 over PPP / PPPoE / VLAN / Ethernet

The E320 router supports a maximum of 8000 dual-stack subscribers over PPP in each ES2 4G LM and ES2 10G LM, and 16,000 dual-stack subscribers over PPP in each ES2 10G ADV LM.

Change in existing behavior: New system maximums as described here.

Unified ISSU

- Unified ISSU Support for OSPFv3

You can now perform a unified in-service software upgrade (ISSU) operation in a router while retaining the OSPFv3 configurations in the router. The IP version 6 (IPv6) traffic outage is maintained approximately at 45 seconds during the unified ISSU upgrade phase. You need to perform the following tasks before you initiate the unified ISSU operation:

- Enable the chassis cluster feature in the router to be upgraded.
- Enable the OSPFv3 graceful restart feature in the router to be upgraded.
- Configure all neighboring routers as graceful restart helper routers.
- Set a higher value for the hold timer in the router to be upgraded.

The upgrading router sends graceful restart notifications to neighboring routers during the unified ISSU operation to indicate that it is performing a graceful restart. The graceful restart feature ensures that there is no impact on the network during the upgrade. The router advertises a grace period of 180 seconds in the Type 9 grace link-state advertisement (LSA). The neighboring routers act as helpers during this grace period.

The hold timer value should be sufficient to overcome the internal communication disruptions during the unified ISSU operation. The appropriate value for the timer is displayed in a warning message during the unified ISSU initialization phase.

The existing **overload advertise-high-metric issu** command used to implement the route-around feature is now supported for OSPFv3. This command enables the upgrading router to advertise a maximum cost on all OSPFv3-enabled interfaces when the unified ISSU operation is started, so that the neighboring routers can select an alternative path (that is, the neighboring routers can route around the upgrading router).

Change in existing behavior: Existing feature extended as described here. In lower-numbered releases, ISSU support was provided only for OSPFv2.

Unsupported Features

The JunosE Release 14.2.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

- 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 the following error message:

% feature not supported on this line card

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 and IPv6
- 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
4.2.0	Not applicable	12.2.1, 12.3.0, 13.0.0
4.3.0	Not applicable	13.0.0, 13.1.0, 13.2.0
4.4.0	Not applicable	13.2.0, 13.3.0, 14.1.0

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 or authorization requests to a RADIUS server, AAA internal limits prevent the actual number of outstanding authentication or authorization requests from exceeding 9600. These internal AAA limits apply only to authentication or 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 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

- When you specify the **show interfaces** *interfaceType interfaceSpecifier* command, if you not leave a space between the *interfaceType* and the *interfaceSpecifier* options, the command processes correctly and displays the settings configured on the corresponding interface. For example, if you enter the **show interfaces gi2/0/0** command, which denotes a Gigabit Ethernet interface on slot 2, adapter 0, and port 0, the system validates this command accurately and considers this command to be the same as **show interfaces gi 2/0/0**. Similarly, if you specify the **show interfaces l2tp members** command, the system analyzes the “l2tp” string to denote a LAG interface with a name of “2tp”.

This method of processing of the interface names occurs because the system processes the letters in the interface name until it encounters an integer. The string of characters in the interface name is mapped to a valid interface type, if a match exists, and the remainder of the name is treated as the interface specifier. This behavior is expected when you specify the interface type and specifier options in the **show interfaces** command without a space separating the two options.

- 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	

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

Deprecated Command	Command Mode	Preferred Command
policy-list	Global Configuration	ip policy-list
radius disconnect client	Global Configuration The RADIUS Disconnect Configuration mode has been removed from the CLI.	subscriber disconnect
rate-limit-profile	Policy List Configuration	rate-limit-profile in Classifier Group Configuration mode
remote-loopback	Controller Configuration	
router-name	Domain Map Configuration Still available in Tunnel Group Tunnel Configuration mode	auth-router-name and ip-router-name in Domain Map Configuration mode
show controllers t1/e1	User Exec, Privileged Exec	
show controllers t1 remote	User Exec, Privileged Exec	
show ike certificates	User Exec, Privileged Exec	show ipsec certificates
show ike configuration	User Exec, Privileged Exec	show ipsec ike-configuration
show ike identity	User Exec, Privileged Exec	show ipsec identity
show ike policy-rule	User Exec, Privileged Exec	show ipsec ike-policy-rule
show ike sa	User Exec, Privileged Exec	show ipsec ike-sa
show ip dhcp-external binding	Privileged Exec	show dhcp binding
show ip dhcp-external binding-id	Privileged Exec	show dhcp binding
show ip dhcp-local binding	Privileged Exec	show dhcp binding
show ip dynamic-interface-prefix	Privileged Exec, User Exec	None
show ip mirror interface	Privileged Exec	show secure policy-list
show license l2tp-session	User Exec, Privileged Exec	None
t1 lineCoding	Controller Configuration	None. This command never had any effect.
traffic-class	Policy List Configuration	traffic-class in Classifier Group Configuration mode
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.

- The DHCP external server may not be able to bind all DHCP clients when all of the following conditions exist:
 - The 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 the DHCP external server and either DHCP relay or relay proxy are configured on the same ES2 4G LM.
 - The 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 standalone 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 the DHCP external server and either DHCP relay or relay proxy in the same virtual router.
- Configure the client-facing and server-facing interfaces for the 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 the DHCP external server and either DHCP relay or relay proxy on separate ES2 4G LMs.

Dynamic Interfaces

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

Ethernet

- The hashing algorithm that selects the LAG member link is associated with the IP address of the subscriber client to support QoS. Consequently, a particular flow is always hashed to the same link. When a member link is removed from a LAG bundle, traffic rate is disrupted and traffic flow is reduced. When the link goes down and then comes back up, the traffic flow is automatically redistributed.
- When counting bits per second on a Fast Ethernet or Gigabit Ethernet interface, an E Series router includes 12 bytes for interpacket gap, 7 bytes for preamble, and 1 byte for start frame delimiter, for a total of 20 bytes (160 bits) per packet more than some non-E Series routers. This value therefore shows the total bandwidth utilization on the interface, including both data and overhead.

- To bridge unicast known-DA packets at line rate on both 2-Gbps ports of the GE-2 line module or the GE-HDE module when paired with the GE-2 SFP I/O module, the minimum packet size must be at least 144 bytes.

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

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

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

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

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

- The **show interfaces 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-numbered releases, 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-numbered releases.

- 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.
- 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 an 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]

- 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 nine 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 nine 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 nine or more characters.
- Change VRF descriptions, VRF names, or both so that the combination of associated VRF names and descriptions consists of nine 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.
- 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 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.

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, 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.

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.

LDP

- The LDP database can maintain up to 250 neighbors if you configure the hello and dead intervals (or hold timers) for IGP, such as OSPF or IS-IS, to be higher than their default values. If you set the hello and dead intervals (or hold timers) at their default values, the LDP neighbors start flapping (constantly go up and down) when more than 200 LDP neighbors are present.

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.

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 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.
- 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 about policy configurations 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 svaleClac1
host1(config-policy-list)#filter classifier-group svaleClac1
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 clac15
host1(config-policy-list)#next-hop 1.1.1 classifier-group clac15
WARNING: This rule has replaced a previously configured rule.
host1(config-policy-list)#next-interface atm 1/0.0 classifier-group clac15
WARNING: This rule has replaced a previously configured rule.
host1(config-policy-list)#filter classifier-group clac15
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 the **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 the **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.

- When any of the policy resources near the state of being fully exhausted, any attempt to create, modify, or delete a policy rule or classifier that is associated with a policy already referenced by an interface fails. An error message is displayed stating that the policy resources are exhausted. The policy resource exhaustion trap is also generated because of the resource exhaustion. If you try to create a new policy instead of modifying a policy that is attached to an interface, the policy resources are allocated properly.

This behavior occurs because the policy manager application examines the availability of policy resources in the installed line modules and requires additional resources when you attempt to update or delete a policy referenced by an interface. As a result, the attempt to modify the policy fails with the error message on exhausted policy resources.

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 seven physical ports:

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

SNMP

- SNMP MIBs

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

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

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

Trap Name	Expected Enterprise OID	Enterprise OID Sent by SNMP Agent
juniDapsEventSwitchover	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
juniDapsEventModeMismatch	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
juniDapsEventChannelMismatch	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
juniDapsEventPSBF	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
juniDapsEventFEPLF	.1.3.6.1.4.1.4874.3.2.2.1.2	.1.3.6.1.4.1.4874.3.2.2.1.2.0
juniAddressPoolHighAddrUtil	.1.3.6.1.4.1.4874.2.2.2.1.3	.1.3.6.1.4.1.4874.2.2.2.1.3.0
juniAddressPoolAbatedAddrUtil	.1.3.6.1.4.1.4874.2.2.2.1.3	.1.3.6.1.4.1.4874.2.2.2.1.3.0
juniAddressPoolNoAddresses	.1.3.6.1.4.1.4874.2.2.2.1.3	.1.3.6.1.4.1.4874.2.2.2.1.3.0
juniDhcpLocalServerPoolHighAddrUtil	.1.3.6.1.4.1.4874.2.2.2.2.3	.1.3.6.1.4.1.4874.2.2.2.2.3.0
juniDhcpLocalServerPoolAbatedAddrUtil	.1.3.6.1.4.1.4874.2.2.2.2.3	.1.3.6.1.4.1.4874.2.2.2.2.3.0
juniDhcpLocalServerPoolNoAddresses	.1.3.6.1.4.1.4874.2.2.2.2.3	.1.3.6.1.4.1.4874.2.2.2.2.3.0
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.

- If you perform an SNMP walk or a Get operation of the `igmpInterfaceVersion` object of the IP MIB for an IGMP interface configured as a passive interface (by using the **ip igmp version passive** command), a value of 255 is displayed for this object. You must set the value of this object to be only 1 or 2 using SNMP. If you attempt to set the `igmpInterfaceVersion` object to any value other than 1 or 2 using SNMP, an error message is displayed.
- When you perform an SNMP walk of the `ipAddressTable` object of the IP MIB, the following behavior is observed:
 - The value of the `ipAddressStatus` attribute is always shown as preferred. If the interface is in the administratively down state, the IP address of that interface is not preferred for communication. In such a case, the `ipAddressStatus` attribute is still shown as preferred.
 - The `ipAddressOrigin` object always returns as manual, even when the IP address is not manually configured and a link-local address is used.
 - The `ipAddressStorage` object always returns the default value as volatile(2).
- When you perform an SNMP walk of the `ipv6InterfaceRetransmitTime` and `ipv6InterfaceReachableTime` MIB objects of the IP MIB, the values for these objects depend on the setting you configured using the **ipv6 nd reachable-time** and **ipv6 nd ns-interval** commands in the CLI interface.

The `ipAddressCreated`, `ipAddressLastChanged`, and `ipAddressPrefix` objects of the IP MIB are not implemented in JunosE software. The `ipNetToPhysicalLastUpdate` object might display inconsistent values because it is not fully supported.

- You cannot configure the `ospfv3ExtAreaLsdbLimit` MIB object in the OSPFV3 MIB, which denotes the maximum number of external link-state advertisements (LSAs), on the router because this object is not supported. As a result, the `ospfv3LsdbOverflow` trap is not generated to indicate that the maximum number of external LSAs (limit) in the link-state database has been exceeded. Also, the `ospfv3LsdbApproachingOverflow` trap is not generated to denote that the number of external LSAs has exceeded 90 percent of the limit.

SRC Software and SDX Software

- When you enter the **show sssc statistics delta** command to display the baselined SRC client statistics, the Active IP Interfaces field in the output of the command displays a negative value for the current number of active IP interfaces that the SRC client is aware of. This problem occurs for subscriber policies managed by the SRC client. Also, this scenario occurs if you enter the **baseline sssc** command to set a baseline for the SRC client statistics before viewing the baselined statistics using the **show sssc statistics delta** command. The router implements the baseline by reading and storing the statistics at the time the baseline is set and then subtracting this baseline when you retrieve baseline-relative statistics.

For example, if the number of active IP interfaces at a certain point in time is 4 and if you set the baseline for SRC statistics at that point, the output of the **show sssc statistics delta** command displays the value for the Active IP Interfaces field as 0 (by subtracting current statistical value of 4 from the baseline value of 4). If a subscriber goes down and the number of active IP interfaces reduces to 3, the output of the **show sssc statistics delta** command displays the value for the Active IP Interfaces field as -1 (by subtracting current statistical value of 3 from the baseline value of 4). [Defect ID 193694]

- 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.

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 one of the following happens:

- 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 one of the following happens:

- 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 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 the 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.
- A memory leak is observed on the SRP module when subscriber sessions are flapped in an environment in which 48,000 DHCP proxy client bindings are established. [Defect ID 189488]

Forwarding

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

- When you attach certain hierarchical policies to subinterfaces as input policies, secondary input policies, and output policies, incoming traffic loss can occur when the number of subinterfaces to which the policies are attached exceeds 4600. [Defect ID 86741]
- Ethernet statistics are incorrectly displayed for virtual port 8 of the ES2-S1 GE-8 IOA when that module is paired with the ES2 10G LM or the ES2 10G Uplink LM. [Defect ID 174784]
- On the ES2 10G LM, a VLAN ID of 0 assigned to an interface can prevent packets from being properly forwarded. [Defect ID 176125]
- 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]
- For IP and VLAN policies attached to VLAN subinterfaces on ES2 10G LMs and ES2 10G Uplink LMs, the output policy counters for outgoing control and exception packets are incorrectly displayed in the output of the **show ip interface** and **show vlan subinterface** commands. These counters are not incremented correctly in the VLAN policy output section of the output of the **show vlan subinterface** command and in the IP policy output section of the output of the **show ip interface** command. [Defect ID 190083]

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

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

IS-IS

- 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]
- 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]
- If you configure one subinterface with an IPv6 address and set up IS-IS adjacencies on it, and configure another subinterface with an IPv4 address and enable IS-IS adjacencies on it, the router does not learn the IPv4 routes if the IS-IS metric of the IPv6 interface is lower than the IS-IS metric of the IPv4 interface. Similarly, the router does not learn the IPv6 routes if the IS-IS metric configured for the IPv4 interface is lower than the IS-IS metric of the IPv6 interface. In such a scenario, the output of the **show ip route** and **show ipv6 route** commands indicate that the IPv4 or IPv6 routes are not correctly learned by the router. When an IPv6 subinterface and an IPv4 subinterface contain the same IS-IS default metric value, the router learns only the IPv6 routes and does not learn the IPv4 routes. [Defect ID 191859]

L2TP

- After a unified ISSU completes on a router functioning as an L2TP access concentrator (LAC), traffic outages occur on the L2TP network server (LNS)-facing interface at the LAC in a configuration with 16,000 or 32,000 L2TP sessions over 500 tunnels. [Defect ID 180147]
- If you perform a unified ISSU operation on an E120 router or an E320 router that contains two pairs of line modules configured for stateful line module switchover and functions as an LNS device, the SRP module resets during the unified ISSU process. This problem occurs when any one of the following conditions are met: [Defect ID 186910]
 - A certain number of L2TP subscribers are already connected to the router and more subscriber sessions are attempted to be established during the unified ISSU process.
 - The logged-in L2TP subscribers are logged out and the subscriber sessions are attempted to be reestablished.
 - After the initialization phase of the unified ISSU process is started and completed, a stateful line module switchover is performed and another unified ISSU process is performed while more subscribers are logging in.
- 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 over 16,000 tunnels that are established 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]
- Approximately 25 percent of the total number of L2TP subscriber sessions are terminated and reestablished after a long time (about 25 minutes for 8000 sessions) when an ATM line module on a router that functions as the LAC device is reloaded. [Defect 187515]

MLD

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

Work-around: If an MLDv1 router is present on the network, configure version 1 with the **ipv6 mld-proxy version** command on that network interface. (Version 2 is the default.)

- The default value for the MLDv2 proxy unsolicited report interval timer should be 1 second rather than 10 seconds (the value for v1). [Defect ID 46040]

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.

- 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]

Policy Management

- 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]
- Rate limiters on the ES2 10G ADV LMs might be corrupted when a large-scale update of policy information on the line modules occurs. In such a scenario, the rate limiter on the LMs have a different committed-rate value from the committed-rate value for the rate limiter on the SRP module. As a result, the policy configuration on the ES2 10G ADV LMs and the SRP modules become inconsistent. This problem occurs with approximately 13,000 PPPoE subscribers that are logged out and logged in again. [Defect ID 90738]
- High CPU utilization is seen in the output of the **show utilization** command (a value of 100 is displayed under the cpu (%) column) when you issue the **show qos interface-hierarchy interface { gigabitEthernet | fastEthernet }** command. This problem occurs when you attach QoS profiles to interfaces specified with ATM VC and VLAN queues, bulk VLAN configuration, and QoS parameters. [Defect ID 94097]
- 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.

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.

- 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]
- The **no qos-parameter-define definition** command does not delete the specified QoS parameter definition. [Defect ID 176844]

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

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

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

- Simple shared shaping does not function correctly when it is used for 32,000 subscribers on an ES2 10G ADV LM. However, when you change the shaper to compound shared shaping, it works properly. Also, simple shared shaping does not function correctly for 16,000 subscribers on an ES2 10G ADV LM. [Defect ID 183512]
- When you configure an E120 or E320 router with an ES2 10G ADV LM as a LAC on one side of an L2TP tunnel and as an 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]

SDX Software and SRC Software

- When multiple IPv6 interfaces are configured with policies attached from the SRC software, 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]

Server Card Manager (SCM)

- High availability mode transitions to the pending state when you perform the following steps. The high availability state of the system is displayed in the output of the **show redundancy detail** command.
 1. Configure a shared tunnel-server port on an ES2 4G line module that functions as the primary line module in a redundancy group of line modules.
 2. Bring up a GRE tunnel on the primary line module.
 3. Perform a line module redundancy operation to switch over from the currently active primary to the standby module.

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

Service Manager

- 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]
- 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.

- Activation of service sessions for a subscriber with DHCPv6 over IPv6 bindings using the COA method that uses RADIUS Change-of-Authorization-Request (COA-Request) messages and VSAs does not work if the service session was previously activated using the RADIUS login method that uses Access-Accept messages and VSAs. However, this problem does not occur for IP subscriber service sessions. Also, this problem does not occur if service sessions for subscribers with DHCPv6 over IPv6 bindings are activated only using the COA method. [Defect ID 189403]

Stateful Line Module Switchover (High Availability)

- 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]

Stateful SRP Switchover (High Availability) and IP Tunnels

- A packet loss sometimes occurs during stateful SRP switchover when you use the **ping** command on a router that is configured for OSPF graceful restart, and is connected to a helper router in the OSPF IPv6 broadcast network and another helper router in the OSPF IPv6 backbone area. [Defect ID 181470]
 - ERX7xx model, ERX14xx model, or ERX310 router:
 - > When you use the **ping** command with the IPv6 address of the helper router in the multicast area as the destination address and the loopback address of the helper router in the backbone area as the source address, a packet loss of two 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.

Subscriber Management

- 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.

- 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]

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 5000 ANCP clients resets or an intermediate switch resets.
 - When you reload the line module.

Resolved Known Problems

Release 14.2.0 is based on Release 14.1.0 and incorporates all problem resolutions found in that release. The following problems were reported open in Release 14.1.0 and have been resolved in this release, or have been resolved since the 14.1.0 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.

Policy Management

- 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]
- Stale and incorrect policy parameters are seen on SRP modules when a periodic update of policy settings is performed for subscribers on ES2 10G LMs. This problem occurs during simultaneous logging in and logging out of subscribers. [Defect ID 90802]

SNMP

- When an SNMP walk of the ipifstatsTable and ifTable objects of the IF MIB and of the ipAddressprefixTable object of the IP MIB is performed, incorrect values are returned for these objects. Also, the ipNetTpPhysicalIfIndex object returns an incorrect value and the value of this object in the ipNetToPhysicalTable MIB table is not listed in the ifTable. [Defect ID 94613]

Errata

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

- The *Solution* subsection under the *Troubleshooting PPPoE Interfaces* section in the *JunosE Link Layer Configuration Guide, Chapter 14, Monitoring and Troubleshooting Point-to-Point Protocol over Ethernet* incorrectly contains the following example command on configuring the IP address of a PPPoE interface:

```
host1(config-if)# ip address 164.10.6.71 255.255.255.0
```

This command is not relevant and necessary for configuring packet trace logs for a PPPoE interface.

- The *Bidirectional Forwarding Detection Overview* section in the *JunosE IP Services Configuration Guide, Chapter 5, Configuring BFD* inadvertently contains the following repetitive information:

When configured for protocols like OSPF and IS-IS, BFD employs rapid, periodic, and inexpensive hello messages to detect path activity. You can also configure BFD to function with static routes, combining with the BFD poll bit to detect path activity.

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 54
Physical layer values	<i>Physical and Logical Density Maximums</i> on page 55
Link layer values	<i>Link Layer Maximums</i> on page 58
Routing protocol and performance values	<i>Routing Protocol Maximums</i> on page 63
Policy and QoS values	<i>Policy and QoS Maximums</i> on page 66
Tunneling values	<i>Tunneling Maximums</i> on page 69
Subscriber management values	<i>Subscriber Management Maximums</i> on page 71

System Maximums for E120 and E320 Routers	Section
General router values	<i>General System Maximums</i> on page 74
Physical layer values	<i>Physical and Logical Density Maximums</i> on page 75
Link layer values	<i>Link Layer Maximums</i> on page 77
Routing protocol and performance values	<i>Routing Protocol Maximums</i> on page 82
Policy and QoS values	<i>Policy and QoS Maximums</i> on page 85
Tunneling values	<i>Tunneling Maximums</i> on page 89
Subscriber management values	<i>Subscriber Management Maximums</i> on page 91

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 55.)				
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 55.)	4	–	–	14/24
Gigabit Ethernet ports per chassis (GE-8 I/O modules) (See Note 3 on page 55.)	4/16	–	–	14/96
Gigabit Ethernet ports per chassis (OC3-2 GE APS I/O module) (See Note 4 on page 55.)	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 63 per OC3/STM1	252 63 per OC3/STM1	252 63 per OC3/STM1	252 63 per OC3/STM1
Logical E3s per COCX-F3 line module	12	12	12	12

Table 2: Physical and Logical Density Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Logical fractional EIs (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 58.)	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 58.)	32,768	81,920	96,000	96,000
Ethernet S-VLANs per I/O module				
FE-8 I/O and FE-8 SFP I/O	16,384	16,384	16,384	16,384
GE I/O	16,384	16,384	16,384	16,384
GE-2 SFP I/O	16,384	–	–	16,384
GE-8 I/O	16,384	–	–	16,384
OC3-2 GE APS I/O	16,384	16,384	16,384	16,384

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
Ethernet VLANs per chassis (See Note 2 on page 58.)	32,768	81,920	96,000	96,000
Ethernet VLANs per I/O module (no more than 4096 VLANs per port)				
FE-8 I/O and FE-8 SFP I/O	8192	8192	8192	8192
GE I/O	4096	4096	4096	4096
GE-2 SFP I/O	8192	–	–	8192
GE-8 I/O	16,384	–	–	16,384
OC3-2 GE APS I/O	4096	4096	4096	4096
Ethernet VLAN bulk configuration VLAN ranges per chassis	300	300	300	300
Ethernet VLAN bulk configuration VLAN ranges per line module	300	300	300	300
Ethernet VLAN overriding profile assignments per chassis	200	200	200	200
Ethernet VRRP VRIDs per line module	800	800	800	800
Frame Relay virtual circuits per chassis	2000	5000	12,000	12,000
Frame Relay virtual circuits per line module				
COCX-F3	1000	1000	1000	1000
cOCx/STMx F0	1000	1000	1000	1000
OC48 (ERX1440 router only)	–	–	–	1000
Frame Relay virtual circuits per port				
COCX-F3	1000	1000	1000	1000
cOCx/STMx F0	1000	1000	1000	1000
OC48 (ERX1440 router only)	–	–	–	1000
HDLC interfaces per chassis	4000	10,000	24,000	24,000
HDLC interfaces per line module				
COCX-F3	12	12	12	12
cOCx/STMx F0	2000	2000	2000	2000
CT3/T3 F0	1992	1992	1992	1992
OCx/STMx/DS-3 ATM	8000	8000	8000	8000
OCx/STMx POS	4	4	4	4

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
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 58.)	16,000	32,000	32,000	48,000
PPP interfaces per line module				
COCX-F3	12	12	12	12
cOCx/STMx FO	2000	2000	2000	2000
GE/FE	8000	8000	8000	8000
GE-2	8000	–	–	8000
GE-HDE	8000	–	–	8000
OCx/STMx/DS-3 ATM	8000	8000	8000	8000
OC3/STM-1 GE/FE	8000	8000	8000	8000
OCx/STMx POS	4	4	4	4
OC48 (ERX1440 router only)	–	–	–	1
PPP packet logging				
Aggregate dynamic and static PPP interfaces for which you can log PPP packets per chassis	32	32	32	32
PPPoE service name tables				
PPPoE service name tables per chassis	16	16	16	16
Service name tags per PPPoE service name table (including one empty service name tag)	17	17	17	17
PPPoE subinterfaces				
Subinterfaces per chassis (See Note 1 on page 58.)	16,000	32,000	32,000	48,000

Table 3: Link Layer Maximums Table continued

Feature	ERX310	ERX705 and ERX710	ERX1410	ERX1440
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 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 63.)				
	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 63.)	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 63.)				
Routing table entries (See Note 4 on page 63.)	500,000	500,000	500,000	500,000
ANCP Adjacency Scaling (See Note 5 on page 63.)	5000	5000	5000	5000
BGP-4 peering sessions	1000	1000	1000	1000
BGP for IPv4 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 for IPv4 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 protocol scaling and peering densities				
(See Note 3 on page 63.)				
BGP for IPv6 routes (NLRI)	100,000	1,000,000	100,000	100,000
OSPF for IPv6 routes	25,000	25,000	25,000	25,000
IPv6 routing table entries				
(See Note 3 on page 63.)				
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	8000	8000	8000	8000
(See Note 6 on page 63.)				
Total Martini circuits per chassis	16,000	16,000	16,000	32,767
(See Note 7 on page 63.)				
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	16,384	16,384	16,384	16,384
(See Note 8 on page 63.)				
Outgoing interfaces per chassis	65,536	65,536	65,536	65,536
(See Note 9 on page 63.)				
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	400,000	400,000	400,000	400,000
(See Notes 10 and 11 on page 64.)				
Dynamic extended translations (NAPT) per SM	200,000	200,000	200,000	200,000
(See Notes 10 and 11 on page 64.)				
Response Time Reporter simultaneous operations per VR				
	500	500	500	500
VRRP VRIDs per line module				
	See Ethernet VRRP VRIDs per line module on page 60.			

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 1,024,000.

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 66.)				

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 66.)	1,024,000	1,024,000	1,024,000	1,024,000
Unique hardware policy assignments per line module for modules other than the GE-2, GE-HDE, and OC48/STM16 (See Note 2 on page 66.)	8127/16,255	8127/16,255	8127/16,255	8127/16,255
CAM entries (See Note 3 on page 66.)				
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 12, 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 69.)				
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 69.)				
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 69.)	16,000	16,000	16,000	32,000
L2TP sessions per line module (See Notes 1 and 3 on page 69.)				
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 69.)				
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)	100,000	100,000	100,000	100,000
(See Note 1 on page 71.)				
DHCP local server				
(See Note 2 on page 71.)				
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 71.)	32,000	32,000	32,000	32,000
DHCP relay and relay proxy client (See Notes 1 and 2 on page 71.)				
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 71.)				
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 that 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 that 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 and E320 routers.

General System Maximums

Table 8 lists some general system maximums for the E120 and E320 routers. 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 74.)	3000	3000
Virtual routers and VRFs per line module, combined (See Notes 1 and 2 on page 74.)	3000	3000
ICR Partitions per chassis	640	640
ICR Partitions per line module	64	64

Physical and Logical Density Maximums

Table 9 lists physical and logical density maximums for the E120 and E320 routers. 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 six 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 six active ports and six 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 install the ES2 4G LM only 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 75.)		
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 75.)		
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 75.)		
Gigabit Ethernet ports per chassis (ES2-S3 GE-20 IOA)	120	240
(See Note 2 on page 75.)		
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 and E320 routers. 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 77.)	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 77.)	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 77.)	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 77.)	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 77.)	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 77.)	64,000	96,000
Ethernet S-VLANs per IOA (See Note 6 on page 77.)		
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 77.)	64,000	96,000
Ethernet VLANs per IOA (See Note 7 on page 77.)		
ES2-S1 GE-4 IOA (with ES2 4G LM) (See Note 5 on page 77.)	16,384	16,384
ES2-S1 GE-8 IOA (with ES2 4G LM or ES2 10G LM) (See Note 5 on page 77.)	16,384	16,384
ES2-S1 GE-8 IOA (with ES2 10G ADV LM) (See Note 5 on page 77.)	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 77.)	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 77.)	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 77.)		
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 77.)	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 77.)	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 77.)	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 and E320 routers. 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. Even though 128,001 IP interfaces are supported, only a maximum of 96,000 subscribers are supported per chassis. A combination of single-stack and dual-stack subscribers to use all 128,001 interfaces is not supported.
4. These values are subject to limitations on available SRP module memory, which varies according to your router configuration.
5. 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.
6. 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.
7. 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).
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.

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 82.)	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 Notes 2 and 3 on page 82.)	128,001	128,001
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 4 on page 82.)		
Routing table entries (See Note 5 on page 82.)	500,000	500,000
ANCP Adjacency Scaling (See Note 6 on page 82.)	5000	5000
BGP-4 peering sessions	3000	3000
BGP for IPv4 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 for IPv4 routes	25,000	25,000
IPv6 routing protocol scaling and peering densities		
(See Note 4 on page 82.)		
BGP for IPv6 routes (NLRI)	100,000	100,000
OSPF for IPv6 routes	25,000	25,000
IPv6 routing table entries		
(See Note 4 on page 82.)	100,000	100,000

Table 11: Routing Protocol Maximums Table continued

Feature	E120	E320
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
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 7 on page 82.)	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 8 on page 82.)	16,384	16,384
Outgoing interfaces per chassis (See Note 9 on page 82.)	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 80.	See <i>Ethernet VRRP VRIDs per line module</i> on page 80.

Policy and QoS Maximums

Table 12 lists policy and QoS maximums for the E120 and E320 routers. 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 used only 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 1,024,000.
 5. The number of QoS profiles or policies that you can attach depends on the number of IP (IPv4 and IPv6) interfaces that you have created. However, the maximum number of QoS profiles or policy interface attachments per chassis is limited to 128,000.

Table 12: Policy and QoS Maximums

Feature	E120	E320
QoS queues per line module (See Note 1 on page 85.)	128,000	128,000
QoS profiles configurable per chassis	1000	1000
QoS profile attachments per chassis (See Note 5 on page 85.)	128,000	128,000

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
QoS profile attachments per line module		
ES2 4G LM	16,000	16,000
ES2 10G LM	16,000	16,000
ES2 10G ADV LM	32,000	32,000
QoS scheduler nodes per line module		
	64,000	64,000
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 85.)	1,024,000	1,024,000
Policy egress interface attachments per chassis		
IP interface attachments	128,000	128,000
(See Note 5 on page 85.)		
Policy egress interface attachments per line module		
(See Note 2 on page 85.)		
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

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
ES2 10G Uplink LM VLAN interface attachments	8191	8191
Policy ingress interface attachments per chassis IP interface attachments (See Note 5 on page 85.)	128,000	128,000
Policy ingress interface attachments per line module (See Note 2 on page 85.)		
ES2 4G LM combined IP and IPv6 interface attachments	32,767	32,767
ES2 4G LM combined ATM, GRE, L2TP (LAC only), MPLS, and VLAN interface attachments	16,383	16,383
ES2 10G LM IP interface attachments	16,383	16,383
ES2 10G LM 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

Table 12: Policy and QoS Maximums Table continued

Feature	E120	E320
ES2 10G LM	256,000	256,000
ES2 10G ADV LM	512,000	512,000
ES2 10G Uplink LM	256,000	256,000
Policy statistics blocks (ingress) per line module		
ES2 4G LM	256,000	256,000
ES2 10G LM	256,000	256,000
ES2 10G ADV LM	512,000	512,000
ES2 10G Uplink LM	256,000	256,000
Parent groups (egress) per line module		
ES2 4G LM	49,151	49,151
ES2 10G LM (internal parent groups only)	8191	8191
ES2 10G ADV LM (internal parent groups only)	8191	8191
ES2 10G Uplink LM (internal parent groups only)	8191	8191
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 85.)	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 12, 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 89.)	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 89.)	4000	4000
L2TP sessions per chassis (See Note 2 on page 89.)	60,000	60,000
L2TP sessions per line module with shared tunnel-server ports provisioned (See Note 2 on page 89.)	8000	8000
L2TP sessions per ES2-S1 Service IOA (See Note 2 on page 89.)	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
L2TP tunnels per line module with shared tunnel-server ports provisioned (See Note 2 on page 89.)	8000	8000

Table 13: Tunneling Maximums Table continued

Feature	E120	E320
L2TP tunnels per ES2-S1 Service IOA	16,000	16,000
(See Note 1 and Note 2 on page 89.)		

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)	100,000	100,000
(See Note 1 on page 91.)		
DHCP local server		
(See Note 2 on page 91.)		
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	32,000	32,000
(See Note 3 on page 91.)		
DHCP relay and relay proxy client		
(See Notes 1 and 2 on page 91.)		
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

