

# JUNOS 8.5 Software Release Notes

**Release 8.5R4**  
28 May 2010  
**Part Number: 530-022747-01**  
**Revision R4**

These release notes accompany Release 8.5R4 of the JUNOS software. They describe the documentation for the routing platforms and known problems with the software. JUNOS software runs on all Juniper Networks J-series, M-series, MX-series, and T-series routing platforms.

You can also find these release notes on the Juniper Networks Technical Publications Web page, which is located at <http://www.juniper.net/techpubs/>.

## Contents

Release 8.5 Features .....	5
Hardware .....	5
User Interface and Configuration .....	6
Interfaces and Chassis .....	7
Services Applications .....	11
Routing Protocols .....	13
MPLS Applications .....	15
Multicast .....	15
Routing Policy and Firewall Filters .....	16
VPNs .....	16
Class of Service .....	17
Network Management .....	18
JUNOScope .....	18
System Log .....	19
Changes in Default Behavior and Syntax .....	21
Hardware .....	21
User Interface and Configuration .....	22
Interfaces and Chassis .....	22
Services Applications .....	24
Routing Protocols .....	25
MPLS Applications .....	26

VPNs .....	26
Multicast .....	26
Class of Service .....	26
Network Management .....	27
Current Software Release .....	28
Resolved Issues .....	28
Platform and Infrastructure .....	28
User Interface and Configuration .....	29
Interfaces and Chassis .....	30
Services Applications .....	32
Routing Protocols .....	33
MPLS Applications .....	34
VPNs .....	34
Class of Service .....	34
Routing Policy and Firewall Filters .....	35
Network Management .....	35
Outstanding Issues .....	35
Software Installation .....	35
Platform and Infrastructure .....	36
User Interface and Configuration .....	39
Interfaces and Chassis .....	40
Services Applications .....	44
General Routing .....	46
Routing Protocols .....	46
MPLS Applications .....	47
VPNs .....	48
Class of Service .....	49
Forwarding and Sampling .....	50
Routing Policy and Firewall Filters .....	50
Network Management .....	50
Previous Releases .....	51
8.5R3 .....	51
Software Installation .....	51
Platform and Infrastructure .....	51
User Interface and Configuration .....	53
Interfaces and Chassis .....	54
Services Applications .....	56
General Routing .....	57
Routing Protocols .....	57
MPLS Applications .....	58
VPNs .....	59
Class of Service .....	59
Forwarding and Sampling .....	60
8.5R2 .....	60
Platform and Infrastructure .....	60
User Interface and Configuration .....	61
Interfaces and Chassis .....	61
Services Applications .....	63
Routing Protocols .....	63
MPLS Applications .....	64
VPNs .....	64

Class of Service .....	65
Forwarding and Sampling .....	65
Network Management .....	65
8.4R2 .....	65
Platform and Infrastructure .....	66
User Interface and Configuration .....	66
Interfaces and Chassis .....	67
Services Applications .....	68
Routing Protocols .....	69
MPLS Applications .....	69
VPNs .....	70
Class of Service .....	70
Forwarding and Sampling .....	71
Network Management .....	71
Errata .....	71
Platform and Infrastructure .....	71
User Interface and Configuration .....	77
Interfaces and Chassis .....	78
Services Applications .....	78
General Routing .....	78
Routing Protocols .....	78
MPLS Applications .....	79
VPNs .....	79
Class of Service .....	80
Routing Policy and Firewall Filters .....	80
JUNOS XML API and Scripting .....	81
M-series, MX-series, and T-series Upgrade and Downgrade Instructions .....	81
Upgrade to Release 8.5 .....	81
Downgrade from Release 8.5 .....	84
J-series Upgrade and Downgrade Instructions .....	84
Upgrade and Downgrade Overview .....	85
Upgrade Software Packages .....	85
Recovery Software Packages .....	86
Before You Begin .....	86
Downloading Software Upgrades from Juniper Networks .....	87
Installing Software Upgrades with the J-Web Interface .....	87
Installing Software Upgrades from a Remote Server .....	87
Installing Software Upgrades by Uploading Files .....	88
Installing Software Upgrades with the CLI .....	88
Installing Software Upgrades by Downloading Files .....	89
Installing Software Upgrades from a Remote Server .....	90
Downgrade Instructions .....	90
Downgrading the Software with the J-Web Interface .....	91
Downgrading the Software with the CLI .....	91
Special Instructions for J-series Routers with a 256-MB Compact Flash .....	92
Cleaning Up Files .....	92

Verifying Available Compact Flash Space .....	93
Increasing the Compact Flash Space .....	93
Removing the Swap Partition .....	94
Configuring the Unused Swap Partition .....	94
List of Technical Publications .....	95
Documentation Feedback .....	102
Requesting Technical Support .....	103
Revision History .....	104

## Release 8.5 Features

---

The following features have been added to JUNOS Release 8.5. Following the description is the title of the manual or manuals to consult for further information. For a complete list of manuals, see Table 4 on page 96, Table 5 on page 100, and Table 7 on page 102.



**NOTE:** Juniper Networks will discontinue offering printed documentation for JUNOS software documentation, M-series and T-series hardware installation and PIC guides, and the *JUNOScope User Guide*, starting with JUNOS Release 8.5. The following model numbers will no longer be available:

- JUNOS-DOC-S
- JNCSP-DOC-S
- DOC-M10i-HW-S
- DOC-M120-HW-S
- DOC-M160-HW-S
- DOC-M20-HW-S
- DOC-M320-HW-S
- DOC-M40e-HW-S
- DOC-M7i-HW-S
- DOC-MX960-HW-S
- DOC-T320-HW-S
- DOC-T640-HW-S
- DOC-TX-HW-S

Juniper Networks will continue to include printed Quick Start guides with router shipments, and specific installation documentation will continue to be shipped with field-replaceable units (FRUs).

---

### Hardware

- **T1600 Internet routing node**—The new T1600 routing node has a capacity of up to 800 gigabits per second (Gbps), full duplex (1600 Gbps of any-to-any, nonblocking, half-duplex switching). The routing node provides Gigabit Ethernet, SONET/SDH, and other high-speed interfaces for large networks and network applications, such as those supported by Internet service providers (ISPs).

The T1600 routing node features the following new hardware:

- T1600-SIBs.
- Type 4 FPC (T1600-FPC4) includes two Packet Forwarding Engines. Each Packet Forwarding Engine has a capacity of 50 Gbps throughput.

- Three-input 240-A power supplies.
- T1600 craft interface, which allows you to identify the T1600 routing node.

All other craft interface features are the same as on the T640 craft interface.

The T1600 routing node supports:

- All legacy FPCs and PICs supported on the T640 routing node
- All legacy Routing Engines supported on the T640 routing node: RE-600, RE600, and RE-S-2000
- All other legacy components except the SIBs and two-input 180-A power supplies

You can perform an online upgrade of an operational T640 routing node to a T1600 routing node. The T1600 upgrade kit includes five T1600-SIBs, two three-input 240-A power supplies, and one craft interface. The T1600-FPC4 is not included in the upgrade kit, but can be installed after all components of the upgrade kit are installed and operational.



**NOTE:** T640 routing nodes currently connected to a TX Matrix platform cannot be upgraded to a T1600 routing node.

---

[*T1600 Hardware Guide*]

- **Redundant power supplies**—Two redundant, load-sharing three-input 240-A DC power supplies are supported for the T1600 routing node. Each input consists of 48 VDC and return, each with its own 80-A circuit breaker. The input mode switch on the faceplate allows you to set the DC power supply to either two-input mode (not currently supported), or three-input mode (required for the T1600 routing node). [*T1600 Hardware Guide*]
- **PIC support for M120 router**—The 8-port (Type 3) Gigabit Ethernet IQ2 PIC with SFP is now supported on the M120 router. This PIC is not oversubscribed. However, it supports all the other features of the IQ2 PICs. [*M120 PIC Guide*]
- **IPSec and flow monitoring version 9 for Multiservices PIC**—IPSec and flow monitoring version 9 are now supported on the MultiServices 500 PIC. [*PIC Guide, Services Interfaces, Policy Framework*]

### User Interface and Configuration

- **Improved fragment statistic tracking for multilink bundles on LSQ PICs (M-series and T-series routing platforms)**—This feature improves the statistic tracking for multilink bundles on LSQ PICs. The `show interfaces lsq-bundle-name extensive` command now separately displays self-contained fragments and multipart fragments. [*Interfaces Command Reference*]
- **Support for logical router system administrators**—The primary system administrator can now assign logical router system administrators. Logical router system administrators have limited configuration privileges within a logical router

and can only view the command output for the logical routers to which they are assigned. To configure a logical router administrator, include the `[logical-router logical-router-name]` statement at the `[edit system login class class-name]` hierarchy level. *[System Basics, Feature Guide]*

- **Passive FTP support**—Previously, the JUNOS software supported all FTP usage in active FTP mode only. However, certain FTP servers and firewalls only support passive FTP. The JUNOS software now supports both active and passive mode. *[System Basics]*
- **Automatic reenrollment support for IPSec digital certificates before expiration**—Enables you to configure automatic reenrollment of the current digital certificate before expiration. You can configure the percentage of the validity-end-time (specified in the certificate), when automatic reenrollment should be initiated. This feature is not enabled by default. To configure this feature, include the `auto-re-enrollment` statement at the `[edit security pki]` hierarchy level. *[System Basics, Feature Guide]*

## Interfaces and Chassis

- **System log file transfer enhancements**—Enable the transfer of system log files to defined archive sites. Optionally enables file transfer based on file size, specified transfer time, or specified transfer interval. To configure archive site transfer, include the `archive-sites` statement at the `[edit system syslog file]` hierarchy level. To define the day and time at which you want the file transfer to occur, include the `start-time` statement at the `[edit system syslog file filename archive-sites]` hierarchy level. To define the interval (days, hours, minutes) at which you want file transfers to occur, include the `transfer-interval` statement at the `[edit system syslog file filename archive-sites]` hierarchy level. *[System Basics]*
- **Queue statistics support for logical interface sets**—This feature adds support for displaying queue statistics for a set of logical interfaces. The output for logical interface sets is similar to the output of a single logical interface. To display the set of logical interfaces, enter the `show interfaces interface-set interface-set-name` command. The `terse` option displays the interface set only. The `detail` option displays the aggregate statistics of the interface set. To display queue statistics for a set of logical interfaces, enter the `show interfaces interface-set queue interface-set-name [aggregate | remaining-traffic] forwarding-class class-name` command. The `aggregate` option displays the statistics of the entire set of logical interfaces. The `forwarding-class` option displays the statistics for the forwarding class specified. The `remaining-traffic` option displays only the statistics of the unconfigured logical interfaces that do not have an explicit class-of-service traffic control profile applied. A `clear interface-set` command is supported to clear the statistics of the interface set. *[Interfaces Command Reference]*
- **Nonstop routing support for Bidirectional Forwarding Detection (BFD)**—Routing protocols use BFD for fast liveness detection of their neighbors. With nonstop routing enabled, BFD session state is saved on both the master Routing Engine and the backup Routing Engine. When a Routing Engine switchover event occurs, the BFD session state does not need to be restarted, and peer routers continue to interface with the routing platform as if no change had occurred.

To enable nonstop routing support for BFD, include the following configuration statements:

- [edit chassis redundancy]  
graceful-switchover;
- [edit routing-options]  
non-stop routing;

You can also configure BFD trace options by including the following configuration statements:

- [edit protocols bfd traceoptions flag]  
nsr-synchronization;  
nsr-packet;

To view BFD state replication status, issue the `show bfd session extensive` command. The new `replicated` flag appears in the output for this command when a BFD session has been replicated to the backup Routing Engine. [*High Availability*]

- **Graceful Routing Engine switchover for extended DHCP relay agent (M-series, MX-series, and T-series routing platforms)**—The extended DHCP relay agent supports graceful Routing Engine switchover (GRES) on all routing platforms that contain dual Routing Engines. To support GRES, the DHCP relay agent automatically mirrors (replicates) information about the state of bound DHCP clients from the master Routing Engine to the backup Routing Engine.

GRES support for the DHCP relay agent prevents the loss of state information for active DHCP clients. If mastership switches to the backup Routing Engine on routing platforms that contain dual Routing Engines, the secondary DHCP relay agent waits in a warm state and resumes control from the failed primary DHCP relay agent. The new primary DHCP relay agent restores state information about DHCP clients that were active before the switchover to the new master Routing Engine.

To enable GRES support for the extended DHCP relay agent, include the `graceful-switchover` statement at the [edit chassis redundancy] hierarchy level. You cannot disable graceful Routing Engine switchover support for the extended DHCP relay agent when the entire router is configured for graceful Routing Engine switchover. [*High Availability*]

- **FreeBSD upgrade**—The JUNOS software base operating system has been upgraded from FreeBSD version 4.10 to 6.1. FreeBSD, an advanced UNIX operating system, provides the JUNOS software advanced symmetric multiprocessing (SMP) features. When upgrading from JUNOS Release 8.2 or earlier to JUNOS Release 8.5, use the `system software add image no-validate` option. Only use the `install JUNOS software image` when upgrading or downgrading to or from JUNOS Release 8.5. Do not use the `bundle image`. Before upgrading to JUNOS Release 8.5, ensure that the router has a compact flash card with a minimum of 256 MB storage space to avoid disk size restrictions. [*Installation and Upgrade*]
- **Unidirectional link support on 10-Gigabit Ethernet interfaces on MX-series routers Dense Port Concentrators**—Enables 10-Gigabit Ethernet interfaces on

the MX-series DPCs to operate in unidirectional mode. Unidirectional links reduce the number of ports required for broadcast video traffic applications, where most of the traffic flow is in only one direction.

When you configure unidirectional mode on an interface, two additional physical interfaces associated with the original parent interface are created on the same port. The new interfaces function as independently operating links. The transmit-only interface is designated by a `-tx` suffix. The receive-only interface is designated by an `-rx` suffix. You can configure logical interfaces on both of the child interfaces, but not on the parent interface.

To configure unidirectional links, include the `unidirectional` statement at the `[edit interfaces interface-name]` hierarchy level. The configuration can be confirmed using the `show extensive interfaces` command. You can specify encapsulation, MAC address, and MTU size on the TX and RX links. Attributes common to both RX and TX, such as clocking and framing, are configured on the parent interface. [*Network Interfaces and Interfaces Command Reference*]

- **M120 routing platform supported features**—All JUNOS features supported on the M320 routing platform are supported on the M120 routing platform, including the following:
  - VPLS operation without a Tunnel Services PIC
  - VRF table label for aggregated Gigabit Ethernet, 10-Gigabit Ethernet, and VLAN physical interfaces
  - VPLS per-packet load balancing in the Packet Forwarding Engine
  - VRF table label with ML-PPP
  - Increase of the aggregated bundle maximum to 16 links
  - For passive monitoring, the dynamic flow capture (DFC) MIB
  - PPPoE encapsulation on the ATM2 PIC

[*Network Interfaces*]

- **PPP over Ethernet server on M120 routers**—This feature adds support for a Point-to-Point Protocol (PPP) over Ethernet (PPPoE) server on the IQ2 PIC of an M120 router serving as an access concentrator.

The PPPoE server is implemented per Ethernet logical interface. The interface can support VLAN tags but does not support stacked VLANs. Multiple PPPoE sessions can be established on the interface to multiple PPPoE clients.

To configure a PPPoE server on an M120 Ethernet logical interface, first specify the `ppp-over-ether` encapsulation option at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level. Then, include the `pp0` statement for the pseudo PPPoE physical interface at the `[edit interfaces]` hierarchy level. Finally, include the `server` statement at the `[edit interfaces pp0.0 unit logical-unit-number pppoe-options]` hierarchy level. [*Network Interfaces, Interfaces Command Reference*]

- **Separate alarm category for T1 CRC errors**—This feature provides support for a separate alarm category for T1 cyclic redundancy check (CRC) errors. You can

configure a major error rate alarm threshold and a minor error rate alarm threshold for CRC errors. The CRC errors are counted by the PIC and can be displayed using the `show interfaces extensive` command. When the threshold is exceeded, an alarm condition is declared. To configure CRC major and minor error thresholds, include the `crc-major-alarm-threshold` and `crc-minor-alarm-threshold` statement at the `[edit interfaces t1-interface t1-options]` hierarchy level. Specify the error rate as the number of errors per number of bits. For example `1e-3` is one error in  $10^3$  superframes and `5e-6` is 5 errors in  $10^6$  bits. This feature is supported on the Channelized OC3 IQ PIC, Channelized OC12 IQ PIC, and Channelized T3 IQ PIC PICs. It is supported only on T1 interfaces that are part of channelized OC3, OC12, and DS3. This feature is not supported on the lower-speed T1 interface. *[Network Interfaces, Interfaces Command Reference]*

- **Configurable timeout to clear the PPP loop-detected flag (M-series and T-series routers)**—This feature provides a configurable timeout to reset the loop detected flag. When a Point-to-Point Protocol (PPP) session detects a loop, the loop detected flag is set. If the flag is not cleared by the protocol after the loopback is cleared, this timer will clear the flag after the specified time has elapsed. To configure the clear PPP loop detected flag timer, include the `loopback-clear-timer` statement at the `[edit interfaces interface-name unit logical-unit-number pap-options]` hierarchy level and specify the number of seconds to wait. *[Network Interfaces]*
- **Cisco-compatible Frame Relay encapsulation (M120 and M320 routers)**—This feature supports Cisco-compatible Frame Relay encapsulation. Support on the M320 router requires an enhanced FPC. To configure Frame Relay encapsulation on a physical interface, include the `encapsulation` statement at the `[edit interfaces interface-name]` hierarchy level and specify the `frame-relay-ether-type`, `frame-relay-ether-type-tcc`, or `extended-frame-relay-ether-type-tcc` encapsulation type. To configure Frame Relay encapsulation on a logical interface, include the `encapsulation` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level and specify the `frame-relay-ether-type` or `frame-relay-ether-type-tcc` encapsulation type. *[Network Interfaces]*
- **Passive flow monitoring support for the MX960 router**—This feature adds support for passive flow monitoring on the MX960 router. It provides the capability to set passive monitor mode on the Gigabit and 10-Gigabit Ethernet interfaces and to perform MPLS label popping. No new configuration statements are introduced. All needed configuration commands are currently documented. *[Network Interfaces, MPLS Applications, Policy Framework]*
- **Eight-queue support for legacy 2-port DS3 PICs**—Adds eight-queue capability to the `max-queues-per-interface` statement configured at the `[edit chassis]` hierarchy level. To configure a total of eight queues, use two interfaces on port 0 and port 2, respectively, since the Mq chip cannot be configured on contiguous streams. Port 1 and port 3 remain unused. This configuration enables the user to configure eight queues on the following PICs: the Quad T3 (i2c: 0x0206) and the Quad E3 (i2c: 0x0207). *[System Basics]*

## Services Applications

- **Stateful GRES on services PICs**—Graceful Routing Engine switchover (GRES) is now supported for IPsec services on AS and MultiServices PICs with preservation of service state. [*Services Interfaces, High Availability*]
- **L2TP support on MultiServices 400 PIC**—Layer 2 Tunneling Protocol services are now supported on M120 routers on MultiServices 400 PICs, in addition to AS PICs and MultiServices 100 PICs. [*Services Interfaces, PIC Guides*]
- **Support for NAT pools in the packet gateway and packet gateway controller**—Enables you to configure separate NAT pools for the packet gateway and the packet gateway controller. To configure the NAT pool to be remotely controlled by the packet gateway controller, include the `remotely-controlled` statement at the new `[edit services nat pool hierarchy pgcp]` hierarchy level. In addition, you can now configure the number of ports required for video and voice flows within a NAT pool by including the `ports-per-session` statement at the `[edit services nat pool hierarchy pgcp]` hierarchy level. The `ports-per-session` statement has been deprecated from the `[edit services pgcp]` hierarchy level. The `show services nat pool pgcp` command has been added and enables you to display information about the number of ports configured for the NAT pool and whether the pool is remotely controlled by the packet gateway controller. [*Services Interfaces*]
- **Enhanced packet gateway monitoring information**—Enables you to obtain enhanced monitoring information for the packet gateway. The `show services pgcp gate` command has been deprecated and is replaced by an enhanced `show services pgcp gates` command. You can now display information in `brief`, `extensive`, and `count` formats. You can also display information about specific gate IDs. The `show services pgcp termination` command has been deprecated and is replaced by an enhanced `show services pgcp terminations` command. You can now display information in `brief`, `h248`, or `count` formats. You can also display information based on the termination prefix. The `show services pgcp root-termination` command has been added and displays information about H.248 root termination. [*System Basics and Services Command Reference*]
- **Support for multiple virtual packet gateways**—You can configure up to eight virtual packet gateways on a JUNOS routing platform. Each packet gateway is associated with a different PIC, and provides the full set of packet gateway functions. Customers who need a high granularity of service types can create multiple packet gateways, each with its own CoS settings. Customers who need to scale their infrastructure beyond the capacity of a single service PIC can achieve this scalability with multiple packet gateways. [*System Basics and Services Reference, Multiplay Solutions Guide*]
- **Support for bidirectional NAT (J-series Services Routers only)**—Enables you to specify the Network Address Translation (NAT) type for a particular term as either traditional (symmetric) NAT or full-cone NAT, in which all requests from the same internal IP address are mapped to the same external IP address. Port mapping is also supported, but not required. To configure this feature, include the `nat-type` statement at the `[edit services nat rule rule-name term term-name]` hierarchy level. By default, symmetric NAT is supported. [*Services Interfaces*]
- **Extended support for Packet Gateway Control Protocol**—The JUNOS software now supports the following H.248v3 functionality as defined in the *Gateway control protocol v3, ITU T Recommendation H.248.1*, September 2005:

- Support for the H.248 segmentation package
- Support for stream-level statistics
- Support for the context-ID list
- Support for the request-ID for signals
- Improved alignment with the ServiceChange procedures in Annex F of the ITU recommendation

[*Services Interfaces*]

- **Support for acknowledging transaction responses (three-way handshake)**—The packet gateway software supports the acknowledgment of transaction responses, which optimizes the transaction responder's resources. By acknowledging the receipt of a transaction response, the transaction initiator states that it will not retransmit the transaction, which allows the transaction responder to discard information about the transaction response. [*Multiplay Solutions Guide*]
- **Support for transmission and retransmission timers on the packet gateway**—The packet gateway software supports a number of properties and algorithms to determine when it transmits and retransmits PGCP transaction requests and responses. These timers are compliant with Annex D of the *Gateway control protocol v3, ITU-T Recommendation H.248.1*, September 2005. [*Multiplay Solutions Guide*]
- **Support for rate limiting in voice traffic**—Enables you to configure rate limiting for voice traffic. Because PGCP flows involve voice traffic, the flows require quality of service that:
  - Provides the bandwidth that the flow requires.
  - Ensures that flows do not consume more resources than they need.
  - Regulates flows that are nonconforming and present vastly greater rates of traffic.

This quality of service is provided through a two-rate three-color policing functionality on the MultiServices PIC. This policer complies with RFC 2698, *A Two Rate Three Color Marker*, September, 1999. With the rate-limiting capability, the MultiServices PIC can police flows to conform to:

- Committed information rate (CIR)
- Peak information rate (PIR)
- Committed burst size (CBS)
- Peak burst size (PBS)

You use rate limiting with PGCP gates. To enable rate limiting for a PGCP gate, you need to provide traffic management package (TMAN) parameters in the PGCP signaling commands that operate on gates. You configure the following parameters on the PGC:

- Tman/sdr—Sustained data rate. This parameter provides the CIR.

- Tman/pdr—Peak data rate. This parameter provides the PIR.
- Tman/mbs—Maximum burst size. This parameter provides the burst size. Both the CBS and the PBS defined in RFC 2698 map to the maximum burst size.

To view rate-limiting statistics on the packet gateway, use the `show services pgcp gates gateway-name extensive` command or the `show services pgcp gates gateway-name gate-id gate-id extensive` command. [*Services Interfaces*]

- **Support for IPv6 addresses in adaptive services**—Enables you to configure IPv6 addresses for adaptive services such as network address translation (NAT), stateful firewall, intrusion detection service (IDS), and class of service (CoS). To configure a rule for NAT, stateful firewall, IDS, and CoS, you can now type an IPv6 address or prefix for destination prefixes, source prefixes, source addresses, destination addresses, source address ranges, and destination address ranges. In the current release, IPv6 addresses are not supported when application layer gateways (ALGs) are configured in the same rule.

To configure a NAT pool, you can now include an IPv6 address or prefix in the `address` and `address-range` statements at the `[edit services nat pool]` hierarchy level. This NAT pool can also be used by the packet gateway service.

To configure the router to aggregate IPv6 traffic before passing the events to IDS processing, you can include the `destination-prefix-ipv6` and `source-prefix-ipv6` statements at the `[edit services ids rule rule-name term term-name then aggregate]` hierarchy level.

You can now view IPv6 addresses for NAT pools using the `show services nat pool` command. [*Services Interfaces; System Basics and Services Command Reference*]



**NOTE:** IPv6 addresses are not supported when the `gateway-address` statement is configured at the `[edit services pgcp]` hierarchy level or at the `controller-address` statement in the `[edit services pgcp gateway gateway-name gateway-controller gateway-controller-name]` hierarchy level.

---

## Routing Protocols

- **Bidirectional Forwarding Detection (BFD) hold-down timer (M-series, MX-series, and T-series platforms)**—Enables you to configure a BFD hold-down interval for static routes and EBGP peers to specify how long a session must remain up before a state change notification is sent. To configure the hold-down interval, include the `holddown-interval milliseconds` statement at the `[edit routing-options static route address bfd-liveness-detection]` or the `[edit protocols bgp neighbor address bfd-liveness-detection]` hierarchy levels. No other protocols are currently supported for this feature. Use the `show bfd extensive` operational mode command to verify your configuration. [*Routing Protocols, Routing Protocols and Policies Command Reference*]
- **BGP remote next-hop support for single-hop EBGP peers**—Enables you to specify that for single-hop External Border Gateway Protocol, a peer should

accept a remote next hop with which it does not share a common subnet. Previously, next-hop routes that did not share a common subnet were discarded. Both peers must support BGP route refresh. You can configure this feature at the global, group, or neighbor level. To configure this feature, include the `accept-remote-nexthop` statement at the [edit protocols bgp], [edit protocols bgp group *group-name*], or [edit protocols neighbor *address*] hierarchy levels. The statement is also supported for logical routers and the VPN routing and forwarding (VRF) routing-instance type. You must also configure an import routing policy for the EBGp peer to specify the remote next-hop address. Do not configure the `multihop` statement at the same time. [Routing Protocols]

- **IPv4 source routing support disabled by default (M-series, MX-series, and T-series platforms)**—IPv4 source routing is now disabled by default. To enable source routing on IPv4, include the `ip-statement` at the [edit routing-options source-routing] hierarchy level. We recommend that you not enable IPv4 source routing. [Routing Protocols]
- **IS-IS hold-down timer for subsequent SPF calculations**—Enables you to configure a time to hold down, or wait, before running subsequent shortest path first (SPF) calculations after a specified maximum number of calculations have occurred in succession. The default number of SPF calculations that can run before the hold-down timer begins is three. You can configure a range from one through five SPF calculations to run in succession after a network topology change is detected. To configure the time to hold down subsequent SPF calculations, include the `holddown milliseconds` statement at the [edit protocols isis spf-options] hierarchy level. To specify a maximum number of SPF calculations to run in succession before the hold-down timer begins, include the `rapid-runs` statement at the [edit protocols isis spf-options] hierarchy level. In addition, you now use the `delay milliseconds` statement to configure the time to delay running an SPF calculation after a network topology change is detected. Include the `delay milliseconds` statement at the [edit protocols isis spf-options] hierarchy level. The `spf-delay` statement, which you previously used to configure the SPF algorithm delay, has been deprecated. Use the new `show isis overview` operational mode command to verify your configuration. [Routing Protocols Configuration, Routing Protocols and Policies Command Reference]
- **OSPF hold-down timer for subsequent SPF calculations**—Enables you to configure a time to hold down, or wait, before running subsequent shortest path first (SPF) calculations after a specified maximum number of calculations have occurred in succession. The default number of SPF calculations that can run before the hold-down timer begins is three. You can configure a range from one through five SPF calculations to occur in succession after a network topology change is detected. To configure the time to hold down subsequent SPF calculations, include the `holddown milliseconds` statement at the [edit protocols (ospf | ospf3) spf-options] hierarchy level. To specify a maximum number of SPF calculations to run in succession before the hold-down timer begins, include the `rapid-runs` statement at the [edit protocols (ospf | ospf3) spf-options] hierarchy level. In addition, you now configure the SPF algorithm delay with the `delay milliseconds` statement at the [edit protocols (ospf | ospf3) spf-options] hierarchy level. The `spf-delay` statement, which you previously used to configure the time to delay running an SPF calculation following a network topology change, has been deprecated. Use the `show ospf overview` operational mode command to

verify your configuration. [*Routing Protocols Configuration, Routing Protocols and Policies Command Reference*]

- **OSPF graceful restart enhancement**—Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helper router. You can now disable strict LSA checking by including the `no-strict-lsa-checking` statement at the `[edit protocols (ospf | ospf3) graceful-restart]` hierarchy level. [*High Availability*]

## MPLS Applications

- **RSVP for unnumbered interfaces**—The JUNOS software supports RSVP traffic engineering over unnumbered interfaces. Traffic engineering information about unnumbered links is carried in the IGP traffic engineering extensions for OSPF, as described in RFC 4203 (OSPF extensions). In addition, the JUNOS software now supports unnumbered extensions for IS-IS, as described in RFC 4205 (IS-IS extensions). Unnumbered links can also be specified in MPLS traffic engineering signaling, as described in RFC 3477. This feature enables you to avoid having to configure IP addresses for each interface participating in the RSVP-signaled network. To configure RSVP for unnumbered interfaces, you must configure the router with a router ID using the `router-id` statement specified at the `[edit routing-options]` hierarchy level. The router ID must be available for routing (you can typically use the loopback address). The RSVP control messages for the unnumbered links are sent using the router ID address (rather than a randomly selected address). To configure link protection and fast reroute on a router with unnumbered interfaces enabled, you must configure at least two addresses. In addition to the router ID, we recommend that you configure a secondary interface on the loopback. [*MPLS Applications*]
- **LSP teardown for BFD**—Enables you to configure an RSVP LSP to be torn down and resigaled when its BFD session goes down. Traffic can be switched to a standby LSP if one is configured and available. Any actions performed are logged. To configure LSP teardown, include the `failure-action teardown` option for the `bfd-liveness-detection` statement at the `[edit protocols mpls label-switched-path lsp-name oam]` hierarchy level. [*MPLS Applications*]

## Multicast

- **IGMP join and leave recording**—Enables the recording of certain IGMP join and leave events on the entire routing system or on specified IGMP interfaces. To configure accounting on the entire routing system, include the `accounting` statement at the `[edit logical-routers logical-router-name protocols igmp]` or `[edit protocols igmp]` hierarchy level. To enable or disable accounting on individual interfaces, include the `accounting` or `no-accounting` statement at the `[edit logical-routers logical-router-name protocols igmp interface interface-name]` or `[edit protocols igmp interface interface-name]` hierarchy level. [*Multicast*]
- **IGMP snooping for MX-series routers**—Enables you to configure general multicast snooping parameters and specific IGMP snooping parameters to allow Layer 2 interfaces to “snoop” Layer 3 protocols for multicast-related information. To configure IGMP snooping, include the `multicast-snooping-options` and `igmp-snooping` option statements at the `[edit routing-instances]` hierarchy level.



**NOTE:** IGMP snooping is not supported on integrated routing and bridging (IRB) interfaces.

---

[Multicast]

### Routing Policy and Firewall Filters

- **Routing policy for source-class and destination-class usage extended to J-series Services Routers**—You can now configure routing policy for source-class usage (SCU) and destination-class usage (DCU) on J-series Services Routers. SCU counts packets sent to customers by performing lookup on the IP source address. DCU counts packets from customers by performing lookup of the IP destination address. Configure the routing policy at the [edit policy-options policy-statement *statement-name*] hierarchy level. Enable packet counting on an interface at the [edit interfaces *interface-name* unit *unit-number* family *family-name* accounting] hierarchy level. [Policy Framework, Network Interfaces]
- **Voice-aware policer support for L2TP on Ethernet IQ2 PICs (M7i, M10i, and M120 routers)**—Enables you to configure a policer on the Layer 2 Tunneling Protocol (L2TP) network server (LNS) that excludes voice calls from the policing bandwidth. You configure the policer at the [edit firewall policer] hierarchy level. [Policy Framework]

### VPNs

- **Point-to-multipoint LSP support for multicast VPNs**—The JUNOS software supports point-to-multipoint label-switched paths (LSPs) for multicast VPNs. Point-to-multipoint LSPs for multicast VPNs are supported for intra-autonomous system (AS) environments (within an AS), but are not supported for inter-AS environments (between ASs). A point-to-multipoint LSP is an RSVP-signaled LSP with a single source and multiple destinations. To configure the properties of the RSVP traffic engineered point-to-multipoint LSP for multicast VPNs for inclusive tunnels, configure the `rsvp-te` statement at the [edit routing-instances *routing-instance-name* provider-tunnel] hierarchy level. To configure the properties of the RSVP traffic engineered point-to-multipoint LSP for multicast VPNs for selective tunnels, configure the `selective` statement at the [edit routing-instances *routing-instance-name* provider-tunnel] hierarchy level. [VPNs]
- **Expanded interface support for the vrf-table-label statement**—Configuration of the `vrf-table-label` statement at the [edit routing-instances *routing-instance-name*] hierarchy level on core-facing interfaces receiving MPLS packets with a null top label has been extended to support the following PICs:
  - 1-port 10-Gigabit Ethernet
  - 1-port 10-Gigabit Ethernet IQ2
  - 10-port Gigabit Ethernet with SFP

[VPNs]

- **Ignore MTU mismatch on Layer 2 circuits (MX-series, M-series and T-series platforms)**—Enables you to configure local interface switching on a Layer 2 circuit to ignore the MTU configuration set for the associated physical interface. This feature enables you to bring up a circuit between two logical interfaces that are defined on physical interfaces with different MTU values. To configure this feature, include the `ignore-mtu-mismatch` statement at the `[edit protocols l2circuit local-switching interface interface-name]` hierarchy level or at the `[edit logical-routers logical-router-name protocols l2circuit local-switching interface interface-name]` hierarchy level. [VPNs]

## Class of Service

- **Hierarchical schedulers (MX-series routers)**—On MX-series routers, you can now apply a class-of-service (CoS) scheduler configuration at one of four different levels. Include the `interface-set` statement to apply the additional level. This statement is supported only on MX-series routers. The supported scheduler hierarchy is as follows:
  - The physical interface (level 1)
  - The interface set (level 2)
  - The logical interface (level 3)

With queueing dense port concentrators (DPCs), users can specify a shaping rate at the physical interface level and a traffic control profile (output traffic control profile) at the logical interface level. This traffic control profile can specify a shaping rate, a guaranteed rate, and a scheduler map. The scheduler map contains the mapping of queues (forwarding classes) to their respective schedulers (parameters for the queue). Queue parameters can specify a transmit rate and buffer management parameters such as buffer size and drop profile.

In metro Ethernet environments, a VLAN typically corresponds to a customer premises equipment (CPE) device, and the VLANs are identified by an inner VLAN tag on Ethernet frames (called the customer VLAN, or C-VLAN, tag). A set of VLANs can be grouped at the DSL access multiplexer (DSLAM) and identified by using the same outer VLAN tag (called the Service VLAN, or S-VLAN, tag). Hierarchical schedulers enable you to provide shaping and scheduling at both levels in addition to the shaping and scheduling that takes place at the physical interface.

To configure hierarchical CoS schedulers, include the following statements at the `[edit class-of-service interfaces]` hierarchy level:

```
interface-sets {
  interface-set-name {
    interface-parameters;
  }
}
```

To apply hierarchical CoS schedulers to a set of interfaces, include the following statements at the `[edit interfaces]` hierarchy level:

```
interface-set ( ifl-set-name | svlan-set | svlannumber ) {
```

```

    ethernet-interface-name {
        interface-parameters;
    }
}

```

[Class of Service, Network Interfaces]

- **Extended support for shaping on aggregated Ethernet interfaces**—You can now configure shaping for aggregated Ethernet interfaces that use interfaces that originate from Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs. No new commands or statements have been added to support this feature.



**NOTE:** You cannot enable shaping on aggregated Ethernet interfaces when there is a mixture of ports from IQ and IQ2 PICs in the same bundle.

---

[Class of Service, Network Interfaces]

## Network Management

- **Support for the client-list, prefix-list, and client-list-name statements in SNMP configuration**—Instead of adding clients individually to an SNMP community, you can now create a list of clients and add this list to the SNMP community. JUNOS software Release 8.5 allows you to create client lists and prefix lists using the `client-list` and `prefix-list` commands. You can include the `client-list-name` statement to add a client list or a prefix list to an SNMP community. Because both the client list and the prefix list are added to an SNMP community when you include the `client-list-name` statement, you must ensure that you do not create a client list and a prefix list with the same name. [Network Management]

## JUNOScope

- **Element management support for the T1600 routing node**—The JUNOScope software allows you to configure and manage the T1600 routing node using element management tools, such as Looking Glass, Configuration Manager, Software Manager, and Inventory Management. The T1600 routing node runs the JUNOS software, which offers many advanced routing and security services. For more information about the T1600 routing node, see the corresponding hardware guide. [JUNOScope]
- **Configuration auditing of partial configuration (configlet)**—JUNOScope users can now audit changes to a part of the configuration running on a router. You can use this feature to compare changes to a part of the running configuration file against a baseline configlet. To audit partial configuration, select **Configuration > Repository > Audit > Audit Partial Configurations**. For more information about auditing partial configurations, see the *JUNOScope 8.5 Software User Guide*. [JUNOScope]
- **Bulk update for imported configuration**—JUNOScope users can now deploy an imported configuration to multiple routers and perform a bulk configuration update. This feature allows you to load an imported configuration or configlet to a group of devices. To load an imported configuration or configlet to multiple routers, select **Configuration > Repository > Load Configuration**. For more

information about loading an imported configuration or configlet to multiple routers, see the *JUNOScope 8.5 Software User Guide*. [JUNOScope]

- **Element management support for MX480 and MX240 Ethernet Services Routers**—The JUNOScope software allows you to configure and manage the MX480 and MX240 Ethernet Services Routers using element management tools, such as Looking Glass, Configuration Manager, Software Manager, and Inventory Management. The MX480 and MX240 Ethernet Services Routers run the JUNOS software, which offers many advanced routing and security services. For more information about the MX480 and MX240 Ethernet Services Routers, see the corresponding hardware guide. [JUNOScope]

## System Log

The following sets of system log messages are new in this release:

- **AUTHD**—Messages generated by the generic authentication service process (authd).
- **FLOW**—Messages generated by the process that handles flows on routers running the JUNOS Enhanced Services software.
- **FPCLOGIN**—Messages generated by the Flexible PIC Concentrator (FPC) login process (pimlogin), which provides direct login access to Physical Interface Modules (PIMs).
- **FWAUTH**—Messages generated by the process that authenticates users when they initiate a connection across a firewall.
- **MCSNOOPD**—Messages generated by the multicast snooping process (mcsnoopd), which enables a Layer 2 device to examine the content of Layer 3 packets to determine which actions to perform.
- **NSD**—Messages generated by the network security process (nsd), which manages firewall configuration on routers running the JUNOS Enhanced Services software.
- **RT**—Messages generated on routers running the JUNOS Enhanced Services software by the Packet Forwarding Engine as it processes packets for security control in real time.
- **RTLOG**—Messages generated on routers running the JUNOS Enhanced Services software by the system log module of the Packet Forwarding Engine as it processes packets for security control in real time.
- **RTLOGD**—Messages generated by the system log utility for real-time processing of packets for security control (rtlogd) on routers running the JUNOS Enhanced Services software.
- **WEBAUTH**—Messages generated by the Web-authentication binary authentication process (webauth) on routers running the JUNOS Enhanced Services software.

The following system log messages are new in this release:

- ASP\_PGCP\_IPC\_MSG\_WRITE\_FAILED
- ASP\_PGCP\_IPC\_PIPE\_WRITE\_FAILED
- CHASSISD\_FASIC\_SRAM\_ERROR

- CHASSISD\_MAC\_ADDRESS\_FABRIC\_ERR
- CHASSISD\_PEM\_NOT\_SUFFICIENT
- DCD\_PARSE\_ERROR\_HIER\_SCHEDULER
- EVENTD\_PIPE\_CREATION\_FAILED
- EVENTD\_ROTATE\_COMMAND\_FAILED
- EVENTD\_ROTATE\_FORK\_EXCEEDED
- EVENTD\_ROTATE\_FORK\_FAILED
- EVENTD\_TRANSFER\_COMMAND\_FAILED
- EVENTD\_TRANSFER\_FORK\_EXCEEDED
- EVENTD\_TRANSFER\_FORK\_FAILED
- LICENSE\_EXPIRED
- LICENSE\_NEARING\_EXPIRY
- VRRPD\_ADVERT\_TIME\_MISMATCH
- VRRPD\_MISSING\_VIP
- VRRPD\_VIP\_COUNT\_MISMATCH

The following system log messages are no longer documented, either because they indicate internal software errors that are not caused by configuration problems or because they are no longer generated. If these messages appear in your log, contact your technical support representative for assistance:

- ASP\_IDS\_NO\_ENTRY
- ASP\_IDS\_NO\_FLOW
- ASP\_IDS\_NO\_FLOW\_OFFSET
- ASP\_IDS\_NO\_MEM
- ASP\_IDS\_NO\_MEM\_FLOW\_OFFSET
- ASP\_IDS\_NO\_PARENT\_FLOW
- ASP\_IDS\_OBJ\_CAC\_FAIL
- ASP\_IDS\_OBJ\_CAC\_FAIL\_OFFSET
- L2TPD\_SESSION\_IFA\_NOT\_FOUND
- LACPD\_DAEMONIZE\_FAILED
- LACPD\_DUPLICATE
- LACPD\_NOT\_ROOT
- LACPD\_PID\_FILE\_LOCK
- LACPD\_PID\_FILE\_UPDATE
- LACPD\_USAGE
- LOGIN\_INVALID\_LOCAL\_USER

- RPD\_KRT\_AFUNSUPRT
- RPD\_KRT\_UNKNOWN\_RTT

The text logged for the ASP\_SFW\_ALG\_PROMOTION\_FAILED tag is now “ALG promotion failed. Stateful firewall application *stateful-firewall-application-name* conflicts with NAT application *nat-application-name* or conflicts with QoS application; request creation of discard flow”. [System Log]

## Changes in Default Behavior and Syntax

---

### Hardware

- **Combinations of PICs**—On Juniper Networks routing platforms, you can typically install any combination of Physical Interface Cards (PICs) in a single Enhanced Flexible PIC Concentrator (FPC). Newer JUNOS services for some PICs can require significant Internet Processor ASIC memory, and some configuration rules might limit certain combinations of PICs on some platforms. To conserve memory, group PICs in the same family together on the same FPC. Ethernet and SONET/SDH PICs typically do not use large amounts of memory. Adaptive Services, Asynchronous Transfer Mode (ATM) 2, Gigabit Ethernet, and IQ serial PICs use more.

Configuration rules might apply to PICs installed on standard Enhanced FPCs on the following routing platforms: M5, M10, M20, M40, M40e, M160, M320, J20, T320, and T640.

Configuration rules do not apply to PICs installed in the following routers or FPCs:

- J-series, M7i, M10i, or M120 routers
- Enhanced Plus FPCs on M-series and J20 routers
- Enhanced Scaling FPCs

When you upgrade the JUNOS software, a warning appears if any configuration rules affect your PIC combinations. If you continue the installation, the PICs appear to be online (the LEDs are on), but the JUNOS software cannot enable them and they cannot pass traffic. As a workaround, you need to plan which PICs to install on the Enhanced FPCs or PIC slots on your routing platform. For specific information about PIC combination rules, consult Technical Bulletin PSN-2007-01-023. Go to <http://www.juniper.net/customers/support> and click **Technical Bulletins**. On the JTAC Technical Bulletins web page, enter PSN-2007-01-023 in the Search field, select the **CS Technical Bulletin ID** radio button, and click **Search**.

- JUNOS Release 8.5 requires a 256-MB or larger compact-flash card for successful installation. The compact-flash size can be obtained from the output of a `show chassis hardware detail` command.

## User Interface and Configuration

- Configuring a blank password using the `encrypted-password` statement at the `[edit system root-authentication]` hierarchy level for root authentication does allow you to commit a configuration, but you will not be able to log in as superuser and get root level access to the router. [*System Basics*]
- **show task replication command**—The new `show task replication` command displays whether or not graceful Routing Engine switchover (GRES) and nonstop routing (NSR) are configured on the routing platform. Output fields are `Stateful replication` and `RE Mode`. The values for these fields can be `Enabled` or `Disabled` and `Master` or `Standby`, respectively. [*System Basics and Services Command Reference*]
- **New DHCP server option (J-series Services Routers)**—Allows you to specify the next server used for DHCP communication after a boot client establishes initial contact. Use of this feature inserts an IP address in to the `siaddr` field of a DHCP protocol packet. To specify the next DHCP server, include the `next-server` statement at the `[edit system services dhcp]`, `[edit system services dhcp pool pool-id]`, or `[edit system services dhcp static-binding mac-address]` hierarchy level. [*J-series Basic Configuration*]
- **show services nat pool detail command**—The `show services nat pool detail` command output now includes an additional counter that registers the number of addresses in use for dynamic source address NAT pools. [*System Basics and Services Command Reference*]

## Interfaces and Chassis

- **VRRP and mixed tagging support**—Mixed tagging (configuring two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing) is supported only for interfaces on Gigabit Ethernet IQ2 and IQ PICs. If you include the `flexible-vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level for a VRRP-enabled interface on a PIC that does not support mixed tagging, VRRP on that interface is disabled. In the output of the `show vrrp` command, the interface status is listed as `Down`. [*High Availability, Interfaces Command Reference*]
- **Output fields for the show system switchover command have changed**—For the `show system switchover` command, the output field `Out of transition` has been changed to `Steady State` and the output field `In transition` has been changed to `Connected`. [*System Basics and Services Command Reference*]
- **M7i router default option**—M7i routers include a default option (`layer-2-3`) at the `[edit chassis fpc number pic number adaptive-services service-package]` hierarchy level that combines the features available in the Layer 2 and Layer 3 service packages. [*Services Interfaces*]
- **Change to the request support information command**—The `request support information` operational mode command now includes output for the `show route summary` command. [*System Basics and Services Command Reference*]
- **Path MTU discovery for GRE and IPIP tunnels**—Provides support for controlling path MTU settings on a per-tunnel basis and on a global basis for GRE and IPIP tunnels. To enable path MTU discovery for all GRE tunnels, include the `gre-path-mtu-discovery` statement at the `[edit system internet-options]` hierarchy

level. To enable path MTU discovery for all IPIP tunnels, include the `ipip-path-mtu-discovery` statement at the [edit system internet-options] hierarchy level. To enable path MTU discovery for a specific tunnel, include the `path-mtu-discovery` statement at the [edit interfaces *interface-name* unit *logical-unit-number* tunnel] hierarchy level.

Path MTU discovery is enabled by default. To disable path MTU discovery for all GRE tunnels, include the `no-gre-path-mtu-discovery` statement at the [edit system internet-options] hierarchy level. To disable path MTU discovery for all IPIP tunnels, include the `no-ipip-path-mtu-discovery` statement at the [edit system internet-options] hierarchy level. To disable path MTU discovery for a specific tunnel, include the `no-path-mtu-discovery` statement at the [edit interfaces *interface-name* unit *logical-unit-number* tunnel] hierarchy level. Path MTU discovery requires that you configure the `tunnel-level` statement and the `system-level` statement simultaneously. [System Basics]

- **Routing instance and the bridge domain identified in show interfaces irb command**—New output fields have been added to display the routing instance name and the bridge domain name in the output of the `show interfaces irb detail` command. [Interfaces Command Reference]
- **MTU change for J-series interfaces**—The MTU has been reduced to 9150 from 9192 on the following J-series interfaces: ADSL2 + PIM, G.SHDSL PIM, Serial, and Dual-Port Serial PIM. [Network Interfaces]
- **show aps command**—The output for the `show aps` operational mode command has been enhanced to include a new field called `admin down for the Int state`. This field indicates that the interface has been administratively disabled using the `disable` statement at the [edit interfaces *interface-name*] hierarchy level. [Network Interfaces Command Reference]
- **show chassis private mac-addresses command**—The `show chassis private mac-addresses` command has been updated so that the output displays a default value of 64 in the `Private count` field. Previously, the default value was 16. [System Basics and Services Command Reference]
- **New account-order statement for L2TP RADIUS accounting**—Include the `accounting-order value` statement at the [edit access profile *profile-name*] hierarchy level. The only accounting-order value currently supported is `radius`. When you configure the accounting order, all users assigned to the profile have RADIUS accounting enabled. You can continue to configure the RADIUS server either at the [edit access] or [edit access profile *profile-name*] hierarchy levels. [System Basics]
- **Digital certificate support**—On J-series Services Routers, the JUNOS software now also supports the Cisco certificate authority (CA) for IPsec profiles to request digital certificates. The JUNOS software continues to support the Microsoft and Entrust CAs for J-series routers. Use the [edit security pki] hierarchy level to configure IPsec. [System Basics]
- **Support for advertisement interval for VRRP IPv6 address**—Include the `inet6-advertise-interval milliseconds` statement at the [edit interfaces *interface-name* unit *number* family inet6 address *destination-prefix* vrrp-inet6-group *group-number*] hierarchy level. The range for the advertisement interval is 100 through 40,950 milliseconds. You can continue to use the `advertise-interval seconds` statement

to configure the advertisement interval for VRRP IPv4 addresses. [*Network Interfaces*]

- **Valid link mode for J-series 4-port Fast Ethernet ePIM**—For this PIM, the only valid value for the link-mode statement at the [edit interfaces fe-fpc/pic/port] hierarchy level is full-duplex. If you specify half-duplex (or full-duplex mode is not autonegotiated), the interface process writes the following message to the system log: "Half-duplex mode not supported on this PIC, forcing Full-duplex mode." [*Network Interfaces*]

## Services Applications

- **New fields for L2TP operational mode command output**—The new Session encapsulation overhead and Session cell overhead fields have been added to the output of the show services l2tp session extensive command. The new fields display overhead information for cells and encapsulation used within an L2TP session. [*System Basics and Services Command Reference*]
- **IP address as the VPG name and the PGC name**—You can now configure an IP address as the virtual packet gateway (VPG) name. However, the IP address is not used in the operation of the VPG. To configure an IP address for the VPG name, include the gateway statement at the [edit services pgcp] hierarchy level.

You can also configure an IP address as the packet gateway controller (PGC) name. However, the IP address is not used for the connection to the PGC. To configure an IP address for the PGC name, include the gateway-controller statement at the [edit services pgcp gateway gateway-name] hierarchy level. [*Services Interfaces*]

- **Gigabit Ethernet interfaces with VLAN tags**—For a Gigabit Ethernet interface that is not configured with a VLAN tag, you can no longer configure the following two statements at the same time: source-address-filter mac-address at the [edit interface interface-name gigeother-options] hierarchy level and the accept-source-mac statement at the [edit interfaces interface-name unit logical-unit-number] hierarchy level. A warning message is displayed that states that you cannot define source filters on the physical interface and on an untagged interface at the same time. [*Network Interfaces*]
- **Support for rsp interfaces on the M120 router**—The M120 router now supports redundancy services PIC (rsp) interfaces. [*Services Interfaces*]
- **request support information command**—The request support information operational mode command now also includes in its output the output for the show krt queue and show krt state operational mode commands. [*System Basics and Services Command Reference*]
- **NAT addresses**—When configuring network address translation (NAT), if you specify the following addresses that do not match the NAT flow or NAT rule, the corresponding traffic is dropped:
  - Addresses specified in the from destination-address statement, when you are using destination translation
  - Addresses specified in the source NAT pool when you are using source translation

[*Services Interfaces*]

- On Ethernet PICs, you can no longer simultaneously configure source filters on physical interfaces and on untagged logical interfaces. [*Network Interfaces*]

## Routing Protocols

- **OSPF traffic engineering link-local identifier**—OSPF traffic engineering now allows you to configure the link-local TE link-state-advertisement packets to include the link-local identifier. Include the `advertise-unnumbered-interfaces` statement at the `[edit protocols ospf traffic-engineering]` hierarchy level. You do not need to include this statement if the Resource Reservation Protocol (RSVP) can signal unnumbered interfaces as defined in RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*. [*Routing Protocols*]
- **Change to clear ospf database command when running nonstop routing**—With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and synchronize the new database with the master Routing Engine. You can also use the `purge` options with any of the options to discard rather than delete the specified LSA entries. [*Routing Protocols Command Reference*]
- **OSPF local route advertisements**—Beginning with JUNOS Release 8.5, OSPF advertises a local route with a prefix length of 32 as a stub link if the loopback interface is configured with a prefix length other than 32. OSPF also continues to advertise the direct route with the configured mask length. [*Routing Protocols*]
- **IPv6 NDP behavior change**—The IPv6 Neighbor Discovery Protocol (NDP) implementation now drops incoming Neighbor Advertisement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461. [*Routing Protocols*]
- **New option for the clear bgp neighbor command**—A new option `soft-minimum-igp` has been added to the BGP operational mode command `clear bgp neighbor`. This command refreshes the outbound state with the MED minimum IGP reset. [*Routing Protocols and Policies Command Reference*]
- **OSPF link-state advertisements**—Beginning with JUNOS Release 8.5, OSPF no longer advertises a router identifier interface that is not configured to run OSPF as a stub network in its link-state advertisements. [*Routing Protocols*]
- **Hold-down interval for multihop EBGP sessions**—If you configure the hold-down interval for a multihop EBGP session, you must also configure a local IP address by including the `local-address` statement at the `[edit protocols bgp group group-name]` hierarchy level. [*Routing Protocols*]
- For BGP communities, when you configure the extended community type using the `community-ids` statement at the `[edit policy-options community name members]` hierarchy level, you must use an AS number for the `src-as` type and an IP address for the `rt-import` type. [*Routing Protocols*]

- The `authentication-type` statement configured at the `[edit protocols ospf area]` hierarchy level has been deprecated. This statement is no longer necessary. [*Routing Protocols*]
- At the `[edit routing-instances]` hierarchy level, you can no longer configure a routing instance with the name "default." [*Routing Protocols*]

### MPLS Applications

- **New option for show ldp neighbor**—The `show ldp neighbor` command has a new option, `neighbor-address`. This option allows you to display LDP neighbor information about a specific LDP neighbor using the LDP neighbor IP address. [*Routing Protocols Command Reference*]

### VPNs

- **Label aggregation for route reflectors in Layer 3 VPNs**—To work with route reflectors in Layer 3 VPN networks, the Juniper Networks M10i router aggregates a set of incoming labels only in these cases:
  - When routes are received from the same peer router
  - When routes have the same site of origin community
  - When routes have the same next hop

The next-hop requirement is important because route reflectors forward routes originated from different BGP peers to another BGP peer without changing the next hop of those routes. [*VPNs*]

- **show route table command**—The output for the `show route table routing-table-name` operational mode command has been enhanced to display the sender and group address for the PMSI: PIM-SM field, which displays information about the PIM sparse mode provider tunnel. [*Routing Protocols Command Reference*]

### Multicast

- **PIM tunnel configuration**—When you configure a Protocol-Independent Multicast (PIM) tunnel in a VPN VRF routing instance but do not also configure PIM in the master instance, an error message is now displayed that indicates this is an invalid configuration. [*Multicast*]

### Class of Service

- **Change to configuration of CoS interfaces**— The JUNOS XML representation of interfaces at the `[edit class-of-service interfaces]` hierarchy level has changed. If you use JUNOScript or NETCONF scripts to configure or retrieve information about interfaces at this level, you must change them to use the new XML tagging. As an example, in JUNOS Release 8.4 and earlier, the following JUNOS XML and X-Path expressions represent two interfaces, `fe-0/0/0` and `ge-0/0/0`:

```

<configuration>
  <class-of-service>      <interfaces>      <name>fe-0/0/0</name>
    <!-- tag elements for other statements -
->    </interfaces>      <interfaces>      <name>ge-0/0/0</name>

    <!-- tag elements for other statements -
->    </interfaces>    </class-of-service> </configuration>
/configuration/class-of-service/interfaces[name='fe-0/0/0']/*
/configuration/class-of-service/interfaces[name='ge-0/0/0']/*

```

The JUNOS XML and X-Path equivalents for JUNOS Release 8.5 and later are as follows:

```

<configuration>  <class-of-service>
  <interfaces>    <interface>      <name>fe-0/0/0</name>

    <!-- tag elements for other statements -
->    </interface>    <interface>
  <name>ge-0/0/0</name>    <!-- tag elements
for other statements - -> </interface>
</interfaces>    </class-of-service> </configuration>
/configuration/class-of-service/interfaces/interface[name='fe-0/0/0']/*
/configuration/class-of-service/interfaces/interface[name='ge-0/0/0']/*

```

The CLI syntax is not changing for CoS interfaces. [*JUNOS XML API Configuration Reference*]

- On the MX-series router, to enable ingress class-of-service (CoS) capabilities for the enhanced queueing Dense Port Concentrator (DPC), you must perform the following steps:
  1. Include the `mode ingress-and-egress` statement at the `[edit chassis fpc number pic number traffic manager]` hierarchy level.
  2. Configure ingress CoS configuration under the `[edit class-of-service]` hierarchy level.

[*Class of Service*]

## Network Management

- For the `igmpInterfaceIfIndex` object in the Internet Group Management Protocol (IGMP) MIB (IGMP-STD-MIB), the routing platform now reports the `snmplIfIndex` value of the logical interface instead of the `ifIndex` value for the index of the table. [*Network Management*]
- `mplsTunnelReoptimized trap`—The `mplsTunnelReoptimized` trap is generated every time the optimization timer expires; that is, when the optimization timer exceeds the value set for the `optimize-timer` statement at the `[edit protocols mpls label-switched-path path-name]` hierarchy level. [*Network Management*]

## Current Software Release

---

The current software release is Release 8.5R4. For information about obtaining the software packages, see “M-series, MX-series, and T-series Upgrade and Downgrade Instructions” on page 81 or “J-series Upgrade and Downgrade Instructions” on page 84, depending on your router platform.

## Resolved Issues

### Platform and Infrastructure

- When a packet's outer label is set to explicit null and the S bit is not set, the LSP ping command does not work. The JUNOS software does not comply with RFC 4182, “Removing a Restriction on the use of MPLS Explicit NULL”. [PR/74963: This issue has been resolved.]
- On M7i and M10i routers, when the system log for the CFEB becomes full, additional messages are discarded instead of overwriting the oldest messages in the log. [PR/79128: This issue has been resolved.]
- When you enable point-to-multipoint LSPs over an outgoing aggregated Ethernet interface that is configured with circuit cross-connect (CCC) switching, the LSP fails to forward traffic and the following error appears in the system log: “nh\_ucast\_add.” As a workaround, disable the interface and LSP, reenable them in that order, and then clear the RSVP session for the LSP. [PR/105884: This issue has been resolved.]
- On M120, M320, and MX960 routers, when you configure override input packet classification, the feature might not work. [PR/271660: This issue has been resolved.]
- On MX-series routers, when unicast RPF is configured on an interface (the `rpf-check` statement is included at the `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy level), the DPC that houses the interface might generate a core file. [PR/275466: This issue has been resolved.]
- If an aggregated Ethernet or aggregated SONET bundle has a large number of aggregate next hops, when a new child link is added or a child interface goes down and comes up, the Packet Forwarding Engine might generate a core file. [PR/276424: This issue has been resolved.]
- If you enable protocol tracing, writes to the hard drive might be blocked and daemons might delay sending packets. This PR only applies to JUNOS software 8.5R2 or higher. [PR/278580: This issue has been resolved.]
- When graceful Routing Engine switchover is configured on a dual Routing Engine system, the backup Routing Engine might generate a core file. [PR/278901: This issue has been resolved.]
- Swapping an IQ2 PIC with a Services PIC in the same PIC slot might cause the router to crash [PR/280505: This issue has been resolved.]
- On M320 and T-series routing platforms, including the `logical-bandwidth-policer` statement at the `[edit firewall policer]` hierarchy level might degrade forwarding performance, cause the Packet Forwarding Engine to generate a core file and stop functioning, or both. [PR/282169: This issue has been resolved.]

- On M10i routers that have Channelized DS3 IQ PICs installed, the Compact Forwarding Engine Board (CFEB) might generate a core file, which also interrupts FPC operation. [PR/283943: This issue has been resolved.]
- When you issue the `request system software add` command and include a file that does not have the `.tgz` extension even though it may be identical to the file with the `.tgz` extension available from the Juniper Networks support Web site, the router reboots. To avoid this problem, use the file names available from the Juniper Networks support Web site. [PR/283948: This issue has been resolved.]
- Under certain circumstances, DHCP discover packets might be leaked to all the configured VRFs. [PR/286139: This issue has been resolved.]
- When a packet larger than the IP MTU size was transmitted, it registered as a microcode error rather than an MTU error. [PR/294485: This issue has been resolved.]
- When you take offline a T640 routing node that has an aggregated Ethernet member link, multicast traffic does not detour to another link. [PR/294732: This issue has been resolved.]
- When an AS or MS PIC are configured as the tunnel interface, IPv6 multicast does not work over IP. The Tunnel PIC does not have this problem. [PR/296352: This issue has been resolved.]
- On routers configured with aggregated SONET or aggregated Ethernet interfaces and multicast next hops, when the aggregated interface flaps, the kernel might restart unexpectedly. [PR/298073: This issue has been resolved.]
- An MPLS frame with an explicit NULL label designated for the Routing Engine might be dropped by the PFE. [PR/298967: This issue has been resolved.]
- A configuration change or Routing Engine switchover might result in a kernel crash when firewalls, CoS, or IPSec are also configured. [PR/300831: This issue has been resolved.]
- On platforms with dual Routing Engines, the Routing Engines might dump core during processing of a BGP UPDATE message with a NEXT\_HOP attribute that is a broadcast address of a local interface. [PR/302236: This issue has been resolved.]

### User Interface and Configuration

- TACACS+ accounting start or stop requests are incompatible with Cisco ACS. The fix is to configure the `no-cmd-attribute-value` statement at the `[edit system tacplus-options]` hierarchy level. When this is enabled, the JUNOS software sets the value of the `cmd` attribute in the TACACS+ accounting start or stop requests to a null string. This is the behavior Cisco ACS expects in order to save accounting requests to the Accounting file; otherwise, the requests are saved to the Administration file. [PR/252472: This issue has been resolved.]
- When two users have a telnet or ssh session on a router, one in configure private mode and the other in configure mode, the telnet session disconnects if the user in configure mode issues the `load patch` command. [PR/274372: This issue has been resolved.]

- Issuing the `show system rollback 1` command results in syslog messages indicating that the router's configuration has been changed by current user. [PR/278392: This issue has been resolved.]
- In JUNOS Release 8.5 and later, an attempt to log in to a router using SSH might fail with a “Could not chdir to home director: No such file or directory” error message. This problem might occur when specific user account configuration is in place and the router is configured to use the TACACS+ server for authentication. The issue arises only if the TACACS+ server has been configured with a `local-user-name` directive that specifies a nonexistent user. [PR/288116: This issue has been resolved.]
- When the filename in the event-script statement is not included at the `[edit event-options policy policy-name then]` hierarchy level, the event policy process (eventd) might generate a core file. [PR/290515: This issue has been resolved.]
- When a configuration group containing a wildcard match for a static route and qualified next hop of a broadcast interface is applied, the routing protocol process (rpd) might exit and dump core. [PR/290712: This issue has been resolved.]
- When you configure the `ip-address` at the `[edit system radius-options attributes nas-ip-address]` hierarchy level in JUNOS Release 8.5 and later, the `nas-ip-address` attribute is not included in the RADIUS packets. [PR/292274: This issue has been resolved.]
- In the J-Web chassis view, the 10-port Channelized E1 IQ PIC is shown with an incorrect interface position, although the interface index is correct. [PR/294957: This issue has been resolved.]

## Interfaces and Chassis

- When you commit firewall and rpf configurations, an erroneous “nh\_jtree\_fe\_prehandler” message might appear on the Packet Forwarding Engine (PFE). This message is informational only and does not indicate an error condition. [PR/96146: This issue has been resolved.]
- On a dual Routing Engine system with graceful Routing Engine switchover (GRES) enabled, when an IPv6 interface is configured with the `loopback` statement at the `[edit interfaces interface-name together-options]` hierarchy level, the backup Routing Engine might report kernel replication errors in the output of the `show system switchover` command. [PR/102164: This issue has been resolved.]
- When you delete or deactivate an interface on a channelized IQ PIC, the PIC might stop operating and generate a core file. [PR/102420: This issue has been resolved.]
- If there is a PIC error and the PIC is coming online again, the system might reset unexpectedly. [PR/241092: This issue has been resolved.]
- When you configure MLPPP or MLFR UNI NNI (FRF.16) bundles on link services IQ interfaces, a certain mix of traffic might cause a lower-priority queue to be starved when packets expire after not being scheduled for some time. [PR/262901: This issue has been resolved.]
- When a Fast Ethernet interface is connected to a Gigabit Ethernet interface that is configured for full duplex without autonegotiation, the information for the Fast Ethernet interface is incorrect in the “Autonegotiation information” section of

the output from the `show interfaces extensive` command. [PR/263957: This issue has been resolved.]

- If a `compression-device` is mistakenly configured under an ATM interface, the JUNOS kernel might dump core, and restart. [PR/265542: This issue has been resolved.]
- Under loaded conditions, the `show interfaces rlsq` command output might display incorrect statistical information because the statistics replies did not arrive in time. [PR/270467: This issue has been resolved.]
- When Routing Engine mastership is repeatedly switched, routing information maintained on the master and backup Routing Engine might be out of sync, causing all Packet Forwarding Engines to reset. [PR/271141: This issue has been resolved.]
- When you issue the `show interfaces extensive` command for an interface to which a Layer 2 input or output policer is applied, the value in the `Dropped frames` field for the policer might be a negative number. [PR/272971: This issue has been resolved.]
- When graceful Routing Engine switchover (GRES) and LSQ (rlsq) interfaces are configured, the `last change` field in the output of the `show interface redundancy` command might be incorrect after a Routing Engine switchover. [PR/273248: This issue has been resolved.]
- On the M320, a signal integrity issue in old clocking hardware might generate inaccurate alarms and errors when the actual clock is working perfectly. This behavior has no operational impact and has been fixed in later releases. [PR/275308: This issue has been resolved.]
- When member links are configured to be part of RLSQ MLPPP bundle, while the RLSQ interface is yet to be configured, error “BAD\_PAGE\_FAULT” is reported by the kernel if “monitoring interface” is executed on this rlsq logical interface. [PR/277689: This issue has been resolved.]
- For a routing node in a routing matrix, when you remove a hardware component from the chassis, alarms are cleared for that component (which is correct). However, alarms are also cleared for all other components of the same type. [PR/278672: This issue has been resolved.]
- If you power off and power on a model RE-A-2000 Routing Engine on a T640 routing node (by issuing the `request system power-off other-routing-engine` and `request system power-on other-routing-engine` commands), the output of the `show chassis hardware` command no longer includes an entry for SPMB 1. [PR/281463: This issue has been resolved.]
- Adding a per-unit-scheduler configuration to a one-port or two-port IQ PIC might cause errors and affect the forwarding state of the ports. [PR/282934: This issue has been resolved.]
- XGE PICs on M120 routers take an unusually long time (up to 1.5 seconds) to send remote-fault messages. [PR/287147: This issue has been resolved.]
- In JUNOS version 8.5 and later, on systems with unnecessary traceoptions enabled, or other configuration that causes high levels of hard drive activity, the Routing Engine might reset with a `watchdog timeout` error. No coredump is generated. As a possible workaround, change the router's configuration to

eliminate unnecessary traceoptions configuration and to minimize other hard disk drive activity. [PR/288011: This issue has been resolved.]

- When you insert an OC192 SONET/SDH PIC that uses XFP optics in to an Enhanced Type 3 FPC on a T640 routing node, the FPC might generate a core file. [PR/288884: This issue has been resolved.]
- Under the following conditions, a logical interface configured for VRRP (the `vrrp-group` statement is included at the `[edit interfaces interface-name unit logical-unit-number family family address address]` hierarchy level) does not initialize properly and the output for it from the `show vrrp summary` command displays the value `bringup` in the `VR State` field: (1) the logical interface is configured with dual VLAN tags (the `vlan-tags` statement is included at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level, (2) the configuration for another logical interface of the same physical interface includes the `vlan-id` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level. The problem can occur even though the interfaces do not belong to the same VRRP group. [PR/288975: This issue has been resolved.]
- When multicast packets are replicated to multiple outbound interfaces at a moderate traffic load, transmit packets might be corrupted. [PR/289353: This issue has been resolved.]
- On MX-series routing platforms, if you take fabric planes offline and the spare planes become active, you might see high traffic drops or continuous high fabric red drops. To recover from continuous high fabric drops, you must switch the fabric planes again. [PR/291541: This issue has been resolved.]
- SONET interfaces which are configured with interface hold up and hold down timers might remain down after an FPC reset or a PIC reset. To restore the interface, (temporarily) remove the interface hold timers. [PR/291707: This issue has been resolved.]

### Services Applications

- When you commit a configuration that does not include either the `pre-shared-key` statement or the `local-certificate` statement at the `[edit security ike policy policy-name]` hierarchy level, the key management process (kmd) generates a core file. [PR/267957: This issue has been resolved.]
- Services PICs (such as the Adaptive Services and MultiServices PICs) do not record correct information in the `SAMPLE-RATE` field in the header of the cflowd packets that they export. [PR/276142: This issue has been resolved.]
- After a routing instance with an `rlsq` bundle is deactivated and activated and then the primary MS PIC is offlined and brought back online, a Routing Engine switchover might result in a kernel database connection error. [PR/292950: This issue has been resolved.]
- The Real-Time Streaming Protocol (RTSP) application-layer gateway (ALG) implementation was not compatible with some RTSP server implementations. [PR/292961: This issue has been resolved.]

## Routing Protocols

- BGP traceoptions incorrectly reports Path Attribute flags with the EXT bit always reset. [PR/51953: This issue has been resolved.]
- If more than 1000 communities are attached to a route, the routing process (rpd) might become unresponsive. You might need to remove the communities and restart the routing process to recover. [PR/77001: This issue has been resolved.]
- The output from the `show route advertising-protocol bgp neighbor-address community community-id` command is not correct if you specify a particular value (such as 11111:2222) for `community-id`. As a workaround, specify the wildcard value `*.*`. [PR/265624: This issue has been resolved.]
- When you activate or deactivate an aggregate route filter (represented by the `aggregate` statement at the `[edit routing-options rib routing-table]` hierarchy level, its contributing members are not reevaluated and the filter continues to function as it did before the change. [PR/270115: This issue has been resolved.]
- The `show ospf route detail` command output displays the optional-capability value for intra-area router routes only. [PR/273809: This issue has been resolved.]
- IS-IS hello packets might not be generated for 8 to 12 seconds during nonstop active routing (NSR) switchover when the Periodic Packet Management process (ppmd) is not pre-programmed ahead of the switchover time. [PR/276823: This issue has been resolved.]
- When both of the following conditions apply, a change in interface status (up or down) causes a BGP status change: (a) there are more than 255 unnumbered interfaces without a destination address (the `unnumbered-address lo0.0` statement is included at the `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy level for more than 255 logical interfaces), (b) the BGP local address (specified by the `local-address` statement at the `[edit protocols bgp group group-name]` hierarchy level) is the last one in the list of addresses included at the `[edit interfaces lo0 unit 0 family inet]` hierarchy level. As a workaround, either use an unnumbered interface that has a destination address or do not set the BGP local address to an unnumbered interface. [PR/277202: This issue has been resolved.]
- When both of the following conditions apply, BGP evaluation of alternate multipaths does not work correctly: (1) an IBGP peer and an EBGP peer both provide the same prefix with the same AS path, (2) the configuration for the EBGP peer includes the `multipath multiple-as` statement at the `[edit protocols bgp group group-name]` hierarchy level but the IBGP peer's configuration does not. [PR/281447: This issue has been resolved.]
- When two Protocol Independent Multicast (PIM) any-source multicast (ASM) routers on a LAN segment have a directly attached receiver to this segment, IIF\_MISMATCH error messages might be displayed preventing the creation of an (S,G) state. This problem occurs with PIM and Multicast VPN configurations. [PR/281662: This issue has been resolved.]
- Using PIM, certain multicast routing topologies might cause delays in multicast route convergence. [PR/282109: This issue has been resolved.]
- Sometimes memory blocks can become corrupted due to an invalid write in to free memory. [PR/283819: This issue has been resolved.]

- A Bidirectional Forwarding Detection (BFD) protocol session might become stuck for different reasons depending on the release in question. For JUNOS Release 8.3 or earlier, the BFD session might be in the failing state. For JUNOS Release 8.3 or later the BFD session might be in init state. As a workaround, issue the `clear bfd session` command to bring the session back up. [PR/286331: This issue has been resolved.]
- When BGP multipath is enabled (the `multipath` statement is included at the `[edit protocols bgp group group-name]` hierarchy level) and route updates arrive from multipath and nonmultipath peers in a certain order, load balancing across paths might stop working correctly. [PR/288694: This issue has been resolved.]
- When BGP deletes a secondary route, the routing process (rpd) might exit unexpectedly and dump core. [PR/290863: This issue has been resolved.]

### MPLS Applications

- When the routing process (rpd) tries to allocate a large number of MPLS labels, it might be restarted incorrectly due to a label space calculation error. [PR/255428: This issue has been resolved.]
- After upgrading to JUNOS Release 8.4 or later, LDP neighborhood could not be established with another vendor's equipment because of a subnet mismatch. The fix adds a new configuration statement, `allow-subnet-mismatch`, that ignores subnet mismatch for the source address in LDP link hello packets. [PR/285933: This issue has been resolved.]
- When the target of the `ping mpls rsvp` command is another vendor's router, the value in the `Local transmit time` field is a UNIX timestamp instead of an NTP timestamp as specified by RFC 4379. [PR/289535: This issue has been resolved.]
- Packet loss can occur following an RSVP auto-bandwidth adjustment. [PR/289553: This issue has been resolved.]
- Other vendor implementations might send status TLV notification messages with the U-bit set to 0 and F-bit set to 1. While such a combination is not recommended according to RFC 5036, the JUNOS software will tear down the LDP session upon receiving such a status TLV message. [PR/290845: This issue has been resolved.]

### VPNs

- When the `tunnel-services` statement is configured at the `[edit routing-instance instance-name protocols vpls]` hierarchy level and a VPLS interface is configured with an MTU, a virtual tunnel interface might flap due to unrelated configuration changes. As a workaround, remove the `tunnel-service` statement in the routing instance configuration. [PR/297141: This issue has been resolved.]

### Class of Service

- When a fragmentation map is applied on a router containing IQ2 PICs, the following error message is displayed: "COS IPC op 25 (FRAGMAP TABLE UPDATE) failed, err 2 (Subtype Unknown)." [PR/239004: This issue has been resolved.]

- On M120 routers, MX-series routers, and on M320 routers with E3-FPCs, MPLS transit traffic with a label stack that performs a pop operation at the penultimate node is not shaped according to the configured transmit rate exact value, which results in more traffic being sent than should be allowed. [PR/282002: This issue has been resolved.]

### Routing Policy and Firewall Filters

- On the MX platform, a firewall filter with the `ip-options` statement included and applied to the loopback interface might not operate correctly. [PR/283215: This issue has been resolved.]

### Network Management

- As a result of a Routing Engine switchover, many processes will be restarted. During this transient stage, Simple Network Management Protocol (SNMP) agent process (snmpd) may generate a syslog message “Header version mismatch & SNMP\_SMS\_HDR\_ERR: problem with hdr size (6) or msg size (0) message in syslog.” This issue is automatically corrected when the switchover process completes, and there is no operational impact afterwards. [PR/77668: This issue has been resolved.]
- A syntax error in the `mib-rfc3811.txt` MIB file prevents SNMP from using the MIB. The file is included in the package accessible at <https://download.juniper.net/software/junos-export/<release>/juniper-mibs-<release>-signed.tgz>, where `<release>` is a JUNOS Release number such as “8.2R4.5”. [PR/80648: This issue has been resolved.]

## Outstanding Issues

### Software Installation

- For hard disks that were originally formatted by JUNOS Release 4.4 or earlier, after you issue the `request system snapshot partition` command, the router cannot boot from the hard disk. As a workaround, issue the `request system snapshot` command before upgrading. [PR/36742]
- If it takes too long to complete an upgrade to the FIPS version of JUNOS, the Routing Engine might restart. [PR/260513]
- When you issue the `request system partition hard-disk` command, the hard disk repartition fails and the disk becomes unusable. The disk can be recovered by taking a snapshot from the compact flash card and rebooting the router. [PR/269493]
- When a hard disk is partitioned, the `/var/empty` directory might not be created. As a result, the router does not accept SSH connections. As a workaround, use the `mkdir` command to create the `/var/empty` directory. [PR/290064]

## Platform and Infrastructure

- When the Monitoring Services PIC is overloaded, the output from the `show services accounting flow-detail` command might freeze. [PR/32896]
- On T-series platforms, a Layer 2 maximum transmission unit (MTU) check is not supported for MPLS packets exiting the routing platform. [PR/46238]
- When you configure a source class usage (SCU) name with an integer (for example, 100) and use this source class as a firewall filter match condition, the class identifier might be misinterpreted as an integer, which might cause the filter to disregard the match. [PR/50247]
- When a Monitoring Services PIC is overloaded with traffic, the FPC might take the PIC offline and repeatedly send the same error message. The error message does not affect normal operation of the FPC and other PICs. As a workaround, restart the FPC and bring the PIC online. [PR/55981]
- Even if you do not configure IPSec, the key management process (kmd) opens UDP port 500. [PR/59054]
- If you configure several DNS servers by including the `name-server` statement at the `[edit system]` hierarchy level, the JUNOS software uses only the first three configured DNS servers. [PR/59172]
- On a Monitoring Services III PIC configured as a dynamic flow capture (DFC) interface (`dfc-fpc/pic/port`), when you configure the DFC interface as the next hop in a forwarding path, port-mirrored packets might become corrupted. [PR/60799]
- In the output of the `show pfe statistics notification` command, the value is incorrect in the field labeled `options or ttl expired (not RE-destined)`. [PR/64951]
- If you configure 11 or more logical interfaces in a single VPLS instance, VPLS statistics might not be reported correctly. [PR/65496]
- If you see warnings like the following: "Warning: Block size restricts cylinders per group to xx." You can safely ignore them. This type of message indicates the maximum number of cylinders per cylinder group as determined by various other parameters. This warning message no longer appears in JUNOS Release 8.5 and later. [PR/65917]
- In a routing matrix configured for graceful Routing Engine switchover (GRES), when the master Routing Engine of a T640 routing node (line-card chassis, or LCC) enters debug mode, it does not release mastership. [PR/66308]
- If you incorrectly configure an aggregate interface, a physical interface does not get added in to the aggregate bundle even if you have corrected the configuration. [PR/69348]
- When a large number of kernel system log messages are generated, the log information might become garbled and the severity level could change. This behavior has no operational impact. [PR/71427]
- On M320 and T-series routing platforms, there is a process that monitors FPCs while they transition to an online state. If an FPC is busy and cannot complete the transition within the time limit, the process might time out and prevent the FPC from coming online. [PR/72364]

- If you configure the same IPv6 address on the `fxp0` interface and another public interface within the same routing instance, the backup Routing Engine might restart. [PR/72573]
- On M320 and T-series routing platforms, when you configure the local gateway of an IPsec tunnel in a routing instance, IPsec might not function properly over a generic routing encapsulation (GRE) tunnel. [PR/73864]
- In the situation where a Link Services (LS) interface to a CE router appears in the VPN routing and forwarding table (VRF table) and if fragmentation is required, Internet Control Message Protocol (ICMP) cannot be forwarded out of the LS interface from a remote PE router that is in the VRF table. As a workaround, include the `vrf-table-label` statement in the configuration. [PR/75361]
- For J-series Services Routers, if you send a real-time performance monitoring (RPM) probe through an IPsec tunnel and the probe includes the `hardware-timestamp` statement at the `[edit services rpm probe owner-name test test-name]` hierarchy level, RPM `icmp-ping` type probes might not work. [PR/75927]
- When you configure the router to log activity with a firewall filter or perform Routing Engine-based sampling, and heavy traffic passes through the router, the following error message might be displayed: “PKTR DMA age error cell counter incremented.” The error indicates that there might be some packet loss in firewall filter logging or Routing Engine-based sampling. However, transit traffic is not affected. [PR/78712]
- On M160 and M40e routers, a hardware error on the Switch Fabric Module (SFM) might cause the board to reboot. [PR/79236]
- When routes in the routing table for a VPLS routing instance go up and down, the count in the `requests to learn an existing route` field of the output from the `show system statistics vpls` command might show a high count (in the tens of thousands) and numerous instances of the following message might be written to the system log: `/kernel: vpls_learn_l2addr(): identical addr and ifl existed: addr <mac-address>, ifl <interface-index>`. There is no operational impact. [PR/80262]
- On the T-series routing platform, when you include the `no-labels` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level, the statement is added to the configuration; however, MPLS labels are still included in the hash key. [PR/80334]
- For Gigabit Ethernet intelligent queuing (IQ) PICs installed in M-series and T-series routing platforms, system log messages for SFP receive power, laser bias, and temperature alarms might alternate between `set` and `clear`. These messages are mostly cosmetic and do not affect performance of the routing platform. [PR/80393]
- If you configure a policer for BGP traffic and a new BGP neighbor is added, it might cause other established BGP sessions to flap. [PR/80599]
- On Fast Ethernet and Gigabit Ethernet PICs, LACP is not supported on an aggregated Ethernet interface that is configured with either `extended-vlan-vpls` encapsulation or `ethernet-vpls` encapsulation. As a workaround, use `vlan-vpls` encapsulation on the aggregated Ethernet interface. This limitation does not apply to aggregated Ethernet interfaces configured on Gigabit Ethernet IQ2 PICs. [PR/94480]

- A firewall filter that matches the forwarding class of incoming packets (that is, includes the `forwarding-class class-name` statement at the `[edit firewall filter filter-name term term-name from]` hierarchy level) might incorrectly discard traffic destined for the Routing Engine. Transit traffic is handled correctly. [PR/97722]
- On J-series Services Routers, you cannot use a USB device that provides U3 features (such as the U3 Titanium device from SanDisk Corporation) as the media device during system boot. You must remove the U3 support before using the device as a boot medium. For the U3 Titanium device, you can use the U3 Launchpad Removal Tool on a Windows-based system to remove the U3 features. The tool is available for download at <http://www.sandisk.com/Retail/Default.aspx?CatID=1415>. (To restore the U3 features, you can use the U3 Launchpad Installer Tool accessible at <http://www.sandisk.com/Retail/Default.aspx?CatID=1411>.) [PR/102645]
- Juniper Networks does not currently support dynamic ARP resolution on Ethernet interfaces that are designated for port mirroring. This causes the Packet Forwarding Engine to drop mirrored packets. As a workaround, you can configure the next-hop address as a static ARP entry by including `arp ip-address` statement at the `[edit interfaces interface-name]` hierarchy level. [PR/237107]
- When a GRE tunnel is configured over multiple physical paths with load-balancing enabled, it might affect GRE keepalive operation and transit traffic. [PR/251652]
- The **IP Option Errors** section in the output from the `show pfe statistics ip options` command does not include counters for all possible types of errors. [PR/254653]
- When you designate a 10-Gigabit Ethernet interface as a link in an aggregated Ethernet bundle (by including the `802.3ad aexstatement` at the `[edit interfaces ge-fpc/pic/port gigether-options]` hierarchy level) and commit the configuration, the operating system might generate a core file and stop operating. [PR/262424]
- On an M20 router, when you include the `route-accounting` statement at the `[edit forwarding-options family inet6]` hierarchy level, the following message might appear in the system log: **Error requesting SET BOOLEAN, illegal setting 32**. The software is functioning correctly. The error can be ignored. [PR/273762]
- In I-chip platforms, if LSI is enabled for an aggregate child physical interface and the child physical interface is not a member the physical interface of multi-physical interface stream (for example, 10x1GE), the child physical interface statistics are double counted. [PR/274396]
- When a GGSN C-PIC sends a packet larger than the MTU of the outgoing interface in a default VRF, ICMP error messages that indicate fragmentation is needed do not reach the C-PIC. [PR/276392]
- Due to a limitation in the Packet Forwarding Engine (PFE), VPN traffic received on a physical interface on an IQ2 PIC might not be counted on the parent aggregated Ethernet physical interface. [PR/284162]
- If a small form-factor pluggable transceiver (SFP) does not respond to a request for diagnostic data, a message is written to the system log. The message is unnecessary because the failure to respond has no operational impact. [PR/293212]
- When a Multilink Point-to-Point Protocol (MLPPP) link is incorrectly added to a Multilink Frame Relay (MLFR) bundle, the kernel crashes. [PR/294885]

- On an M320 router with redundant Routing Engines, when you deactivate an IP address on a 10-Gigabit Ethernet interface and then add a new IP address, the backup Routing Engine might produce a core dump. [PR/297274]
- When CLNS is configured over a logical tunnel interface, the source MAC address gets corrupted. [PR/304323]

### User Interface and Configuration

- On M20 routers, after a Routing Engine mastership switchover, it might not be possible to enter CLI configuration mode on the new master Routing Engine. Also, the `request system reboot` and `request system halt` commands do not clearly fail but do not return the CLI prompt either. [PR/64899]
- In the J-Web configuration editor, when you select **System > Syslog > File > "filename" > Explicit priority**, the J-Web Event Viewer does not show the event ID. When you select **System > Syslog > Time format > milliseconds**, the J-Web Event Viewer does not filter messages. [PR/70523]
- If the configuration includes both commit scripts (at the [edit system scripts commit] hierarchy level) and control characters from the International Organization for Standardization (ISO) C0 set (included at any hierarchy level), an attempt to commit the configuration fails. As a workaround, remove the control characters. [PR/82384]
- Support for the traceoptions log file is provided for the event scripts. [PR/235912]
- The logical router administrator can modify and delete master administrator-only configurations by performing local operations such as issuing the `load override`, `load replace`, and `load update` commands. [PR/238991]
- When an M-series or T-series router is upgraded from JUNOS to JUNOS-FIPS, the `request system snapshot` command does not work. As a workaround issue a `request system snapshot force-fmt` command from the shell. This issue is not present for upgrades from an older version of JUNOS-FIPS to a newer version of JUNOS-FIPS. [PR/252640]
- Even though the `trace` permission is included at the [edit system login class class-namepermissions] hierarchy level, users who belong to the login class receive the following error when they issue the `show log` command: `error: permission denied: log`. As a workaround, add the `trace-admin` permission to the list of permissions. [PR/278950]
- Sometimes, depending on the configuration, key administration might fail to see an MD5 key configured for a BGP peer as part of a group configuration. [PR/283238]
- Use of system log regular expressions to refine the logged messages does not work properly. [PR/295523]

## Interfaces and Chassis

- On aggregated SONET/SDH interfaces, the counter for drops and errors in the `show interfaces` command output does not display the correct value, because the counter does not collect data from the constituent interfaces within the aggregate. [PR/23577]
- On ATM interfaces, when the IP address of a remote device is changed, the output of the `show ilmi interface` command on the local routing platform might continue to display the old IP address for the remote device. [PR/24126]
- If virtual channel identifiers (VCIs) for a large number (approximately 400) of virtual connections (VCs) on an ATM DS3 interface are changed frequently, the interface might mishandle the ATM cells. As a result, OSPF and IS-IS neighbor adjacencies might not remain stable. [PR/25639]
- On a 2-port OC12 ATM2 IQ interface, the total virtual path (VP) downtime might not appear correctly in the `show interfaces` command output. [PR/27128]
- On a 2-port OC12 ATM2 IQ interface, if you configure and then change the virtual path (VP) setting, the SNMP `jnxAtmVpTotalDownTime` counter might be reset. [PR/27131]
- On an OC3 ATM2 intelligent queuing (IQ) interface, when you configure a shaping rate greater than the speed of the OC3 link and commit the configuration, the actual shaping rate might be less than the interface speed. [PR/27459]
- On the ATM2 IQ PIC, when you configure the `atm-l2circuit-mode` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level, the control word sequence number is not reset to 1 after the transmit sequence number reaches 65,535. [PR/31669]
- On M20 and M40 routers, when a physical layer problem affects a SONET/SDH interface, carrier transition statistics might not increment correctly in the output of the `show interfaces extensive` command. [PR/33325]
- When you configure both the bundle link and constituent links at the `[edit logical-routers logical-router-name interfaces]` hierarchy level, the constituent links do not come up. As a workaround, configure the constituent links at the `[edit interfaces]` hierarchy level. [PR/35578]
- On ATM2 DS3 and E3 interfaces, when you configure ATM point-to-multipoint permanent virtual circuits (PVCs), the following error messages might appear in the system log: “/kernel: RT\_COS: COS IPC op 4 (CLASS TO IFL) failed, err 1 (Unknown),” “ssb BCHIP 0: invalid entry type 127 at stream 8 channel 0 for ifl 83,” and “ssb COSMAN: mapping table bind to ifl 83 failed.” There is no operational impact. [PR/36524]
- When an ATM interface configured for circuit cross-connect (CCC) encapsulation receives MPLS packets that exceed 484 bytes, the packets can overflow the buffer and cause the ATM PIC to hang. As a workaround, take the PIC offline and bring it back online. [PR/39918]
- When you apply an IPsec firewall filter to match traffic sent across a generic routing encapsulation (GRE) tunnel and originating from the local routing platform, the local traffic is dropped. Transient traffic is not affected. [PR/44871]

- On a Link Services PIC with Multilink Frame Relay (MLFR) configured, the `ping` command might fail when the data-link connection identifier (DLCI) is greater than 335. [PR/49567]
- On a Link Services PIC, the CLI might incorrectly allow you to configure a logical tunnel interface (interface identifier `lt`); the resulting interface might not work correctly. [PR/49818]
- If an MLPPP LSQ bundle carries a large volume of link fragmentation and interleaving (LFI) traffic and a small proportion of multilink traffic, packets might be dropped on the egress constituent links. [PR/56664]
- For ISDN dialer interfaces in a J-series Services Router, when you configure the `no-keepalives` statement at the `[edit interfaces dlo unit logical-unit-number]` hierarchy level and you issue the `show interfaces dlo` command, the `Link flags` field might still show `Keepalives`. [PR/58520]
- On ISDN interfaces in a J-series Services Router, if you include the `vrf-table-label` statement at the `[edit routing-instances instance-name]` hierarchy level, packets might be dropped from the connection. [PR/59718]
- On ISDN dialer interfaces in a J-series Services Router, if you include the `minimum-links` statement at the `[edit interfaces dlo unit logical-unit-number]` hierarchy level and then deactivate the BRI interface associated with the dialer interface, the output packets counter displayed in the output of the `show interfaces dlo` command might continue to increment. [PR/59986]
- On ISDN dialer interfaces in a J-series Services Router, when you include the `load-threshold 100` statement at the `[edit interfaces dlo unit logical-unit-number dialer-options]` hierarchy level and the 56-Kbps bandwidth threshold is exceeded, the interface does not support additional network traffic and might not activate another BRI interface. [PR/60045]
- If you configure IS-IS, MPLS, and graceful Routing Engine switchover (GRES) and a switchover event occurs, the routing platform might end the PPP IP Control Protocol (IPCP) sessions and renegotiate them if the remote side changed interface MTU settings before the switchover event. [PR/61121]
- If you configure graceful Routing Engine switchover and issue the `request chassis routing-engine master acquire` command, in rare cases the master Routing Engine might fail to relinquish mastership, or the switchover to the backup Routing Engine might take up to 360 seconds. [PR/61821]
- For Automatic Protection Switching (APS) on SONET/SDH interfaces, there are no operational mode commands that display the presence of APS mode mismatches. An APS mode mismatch occurs when one side is configured to use bidirectional mode, and the other side is configured to use unidirectional mode. [PR/65800]
- J4350 and J6350 Services Routers might not have enough data buffers to meet expected delay-bandwidth requirements. Lack of data buffers might degrade CoS performance with smaller-sized packets (500 bytes or less). [PR/73054]

- The JUNOS software does not always correctly handle MTU settings for individual protocol families, as configured by including the `mtu` statement at the `[edit interfaces interface-name unit logical-unit-number family family-name]` hierarchy level. Specifically:
  1. If you explicitly set the MTU to the default value and then remove the `mtu` statement, the `User-MTU` flag in the output from the `show interfaces` command is not removed for the logical interface.
  2. When you remove the `mtu` statement for a nonnegotiable interface, the MTU value is not reset to the default.
  3. When you explicitly set the `mtu` statement to the default value, the `User-MTU` flag might not be set correctly.

[PR/77975]

- If you include the `disable` statement at the `[edit interfaces interface-name]` hierarchy level to disable the ingress interface for a SONET link between two routers that are not configured for APS or other link protection, the egress interface might not be notified. This can cause traffic loss. [PR/78831]
- On the M120 router, for a Forwarding Engine Board (FEB) redundancy group that does not have a primary FEB configured, when a switchover from a nonprimary FEB occurs, the backup FEB does not reboot, and the Flexible PIC Concentrators (FPCs) connected to the previously active FEB remain online. The backup FEB could take minutes to obtain the entire forwarding state from the Routing Engine following a switchover. If you do not want the interfaces to remain online during the switchover for a nonprimary FEB, configure a primary FEB for the redundancy group at the `[edit chassis redundancy feb]` hierarchy level. [PR/80946]
- On J4350 and J6350 Services Routers, if the MTU is set to more than 6 KB for a built-in Gigabit Ethernet port or a 1-port Gigabit Ethernet ePIM, packets might be discarded with an FCS error. [PR/82245]
- If you ping a nonexistent IPv6 address that belongs to the same subnet as an existing point-to-point link, the packet loops between the two point-to-point interfaces until the time to live expires. [PR/94954]
- If the delay between VRRP advertisement packets is set to a small value (such as 100 ms) for a number of VRRP groups, and the router configuration is changed and committed several times in quick succession, the VRRP mastership state might be unstable. In other words, if the value of the `fast-interval` statement at the `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-number]` hierarchy level is 100 for several VRRP groups, and configuration changes are committed several times in quick succession (even changes at other levels of the hierarchy), a VRRP backup router might assume mastership and immediately release it again. As a workaround, set the value of the `fast-interval` statement to 300 or higher. [PR/102111]
- The output of the `show interfaces diagnostics optics` command includes the `Laser rx power low alarm` field even if the transceiver is a type (such as XENPAK) that does not support this alarm. [PR/103444]
- For Gigabit Ethernet interfaces on J-series Services Routers, the `link-mode` and `speed` statements at the `[edit interfaces ge-fpc/pic/port]` hierarchy level are

mutually dependent; that is, if you include one, you must include the other. If you do not, the interface process generates a warning and uses autonegotiated values. For Gigabit Ethernet interfaces on other routing platform types, the `speed` statement is not available, so including the `link-mode` statement alone is valid. Nevertheless, the interface process writes the following message to its log and the system log: **Speed and linkmode duplex settings are mutually required.** (Note further that the `link-mode` statement is actually nonoperational on non-J-series routing platforms, because the only valid value for it is the default, `full-duplex`.) [PR/228857]

- On channelized DS3 interfaces, if you include the `family mlfr-end-to-end` statement at the `[edit interfaces ivinterface-name unit logical-unit-number]` hierarchy level for a logical interface that has a smaller index number than another logical interface for which you include the `encapsulation frame-relay-ppp` statement at the `[edit interfaces ivinterface-name unit logical-unit-number]` hierarchy level, the commit operation fails with the error message **Link encapsulation type is not valid for device type.** As a workaround, configure the indicated `family` statement on a logical interface with a larger index than the logical interface configured with the indicated `encapsulation` statement. [PR/229071]
- When you configure the `default-address-selection` statement at the `[edit system]` hierarchy level, Routing Engine graceful restart may cause GPRS support node (GGSN) services to be unreachable. [PR/232197]
- When you issue a `show chassis ether-switch statistics` command while redundancy is enabled, there is a loss of communication between the two redundant REs for about 2 seconds. [PR/233779]
- On serial interfaces transmitting either 64-byte or 128-byte packets, the effective bandwidth falls when the interface is highly oversubscribed. [PR/235753]
- When a redundant power supply is removed from an M7i or M10i router, the `show chassis environment` command correctly shows the supply's status as **Absent**, but continues to display a temperature for it. [PR/241055]
- On a J-series router, when you upgrade a serial interface to JUNOS Release 8.0 and later, Frame-Relay encapsulation might not work. Frame Relay does work with JUNOS Releases 7.0 through 7.6. [PR/241610]
- When you configure an IPv6 address as the primary or preferred address for an interface (by including the `primary` or `preferred` statement at the `[edit interfaces interface-name unit logical-unit-number family inet6 address ipv6-address]` hierarchy level) and commit the configuration, messages like the following are written to the system log: **DCD\_CONFIG\_WRITE\_FAILED: Interface *interface-name* configuration write failed for an IFA CHANGE: Operation not supported.** [PR/258531]
- On ATM1 PICs, the effective shaping rate is lower than that specified by the values you configure when including the `shaping` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level. As a workaround, set values appropriate for a shaping rate 4.5 percent higher than desired. [PR/268763]
- If you configure the number of minimum links eight, deactivated interfaces are not counted as down and active interfaces are still brought up. [PR/285244]

- The interface hold-timer might not work for channelized subinterfaces. [PR/294654]
- When there was a change in priority or tracking info, the state machine would reset and you would see VRRP go through an idle-backup-master transition. [PR/303701]

### Services Applications

- The output of the `show services nat pool` command displays duplicate entries for a single Network Address Translation (NAT) pool. [PR/34678]
- The `show services accounting flow-detail extensive` command sometimes displays incorrect information about input and output interfaces. [PR/40446]
- When you configure intrusion detection service (IDS) on J-series platforms, including the `threshold` statement at the `[edit services ids rule rule-name term term-name then logging]` hierarchy level has no effect. [PR/46577]
- On Adaptive Services PICs configured for IPSec tunnel redundancy, if there are a large number of tunnels, sometimes a few of the tunnels might switch over to the backup tunnel. [PR/46733]
- On routing platforms configured for Internet Key Exchange (IKE)-based IPSec, if a remote peer using other vendors' equipment does not renegotiate the IKE security association (SA) when it is about to expire and continues to send dead peer detection (DPD) requests on the same SA, the routing platform might not be able to reply to these messages. [PR/47004]
- If the socket buffer becomes full on a remote router, you cannot clear all the IPSec security associations (SAs) from the router. [PR/55189]
- When a routing platform is configured for graceful Routing Engine switchover and Adaptive Services (AS) PIC redundancy, and a switchover to the backup Routing Engine occurs, the redundant services interface (`rsp-`) always activates the primary services interface (`sp-`), even if the secondary interface was active before the switchover. [PR/59070]
- On Monitoring Services I and Monitoring Services II PICs, if the export channel to the external cflowd collector is closed, cflowd records might be lost. As a workaround, restart the PIC. [PR/59432]
- On Monitoring Services II PICs configured for flow collection services, during memory overload conditions, the flow collector interface might create files lacking cflowd records, and these files might not be sent to the external FTP server. [PR/62599]
- When you modify a flow collection configuration and commit the changes, the system log might contain error messages regarding the commit operation. These messages do not affect the operation of the router and can be ignored. [PR/64201]
- On J-series Services Routers, an SNMP query returns a zero value for the data link switching (DLSw) MIB object `dlswTConnTcpConfigKeepAliveInt` even if you implement keepalives. [PR/70002]
- For Adaptive Services II PICs, even if you do not configure flow collector services, a temporary file might be created every 15 minutes in the `/var/log/flowc/`

directory. The file is deleted if there are no clients, and re-created only when a client connects and attempts to write to the file. [PR/75515]

- The destination IP address assigned to a VP interface can be a duplicate of the address assigned to another interface on the router. This can cause issues with forwarding traffic appropriately to the VP interface. [PR/75535]
- On J4350 and J6350 Services Routers, when you insert a Telephony Gateway Module (TGM) 550 PIM and the PIM is in a reset state, the router might not respond to any **show chassis** commands for up to 5 seconds. [PR/78695]
- In BIOS configuration mode, pressing the F10 key to complete a save and exit does not work as expected. The alternative to using the F10 key is to use the **Save and Exit** option from the **Exit** menu. Regardless of which J-series image is loaded on the router, this issue can be seen on the J4350 and J6350 routers with BIOS Version 080011 and on the J2320 and J2350 routers with BIOS Version 080012. [PR/237721]
- The Clear NVRAM option in BIOS configuration mode does not work as expected. Regardless of which J-series image is loaded on the router, this issue can be seen on the J4350 and J6350 routers with BIOS Version 080011 and on the J2320 and J2350 routers with BIOS Version 080012. To help address this issue, you need to note any changes you make to the BIOS configuration. This allows you to revert to the default BIOS configuration when needed. [PR/237722]
- When a packet-gateway subtract command does not include an audit descriptor, an inappropriate error message is returned: ER=444{"An unknown descriptor was received. [PR/240758]
- You might not be able to deactivate and reactivate Packet Gateway Control Protocol (PGCP) services and services state. [PR/253513]
- The gate inactivity duration for CLI values was changed from a default value of 0 to 5, changing the range to 5 to 86,400 (it was 0 to 86,400). A zero value is no longer valid. [PR/253517]
- The JUNOS software incorrectly sends a data inactivity notify message when a termination is **OutOfService**. [PR/254873]
- JUNOS does not reset the T-MAX timer when receiving a provisional response. The T-MAX timer should be reset once a provisional response (pending) is received by the packet gateway controller (PGC). However, the packet gateway (PG) fails to reset and does not send the expected service change with disconnect message to the PGC before the T-MAX or T-Super timers expire. [PR/255360]
- The **tmax-retransmission-delay** statement configured at the **[edit services pgcp gateway gateway-name h248-timers]** hierarchy level does not function correctly. If you configure a value of 60 seconds, the packet gateway (PG) should send the first notify message to the packet gateway controller (PGC) in 60 seconds. If no message is received after 60 seconds, the PG should send a ServiceChange message with method Disconnect. However, after 10 Notify messages are sent by the PG, it sends a ServiceChange message with a duration of 11 seconds (even though it should be 60 seconds). [PR/255386]
- An inactivity notification is sent by the packet gateway (PG) even when inactivity is detected on the Real Time Control Protocol (RTCP) flow. [PR/256115]

- The value in a ServiceChange(Disconnect) message from the router might be 1. It should be the version negotiated between the router and the PGC. This is typically version 3. [PR/256857: This problem is resolved.] [PR/256857]
- When the packet gateway (PG) SIP-TCP (LATCH) Terminations are changed by a Modify message to out of service, the PG should not perform Latch or Relatch operations when the packet is received on the BB or AC side. However, the PG appears to perform the Latch operation and sends a NOTIFY message regarding this operation to the packet gateway controller (PGC). [PR/259356]

## General Routing

- LDP sessions might go down and remain in an inoperative state for a long time (one indication is that the value `OpenSent` or `Closing` persists over time in the `State:` field of the output from the `show ldp session extensive` command). This problem occurs when BGP must evaluate a large number of AS paths as required by the following configuration:
  1. The value of each of several `as-path policy-name` statements at the `[edit policy-options]` hierarchy level is a regular expression containing a large number of AS path index numbers.
  2. Such policies are each specified as the value of a `from as-path` statement at the `[edit policy-options policy-statement statement-name]` hierarchy level.
  3. Several such policy statements are specified as values for the `import` statement at the `[edit protocols bgp]` hierarchy level.

[PR/229273]

- If the `from` clause in a policy refers to the routing table used by a VPN routing and forwarding (VRF) instance, and you change the route distinguisher for that VRF instance, the routes in the routing table become unusable. In terms of configuration statements, the routing table is the value of the `rib` statement at the `[edit policy-options policy-statement policy-name term term-name from]` hierarchy level, and the route distinguisher is defined by the `route-distinguisher` statement at the `[edit routing-instances routing-instance-name]` hierarchy level for the VRF instance. As a workaround, deactivate the `policy-statement` statement temporarily while changing the route distinguisher. [PR/254398]

## Routing Protocols

- When you include the `as-path atomic-aggregate` statement at the `[edit routing-options aggregate defaults as-path]` hierarchy level to manually add the `ATOMIC_AGGREGATE` attribute on a BGP AS path, the attribute is not added. [PR/2527]
- The `metric-out` statement at the `[edit protocols protocol-name group]` hierarchy level incorrectly takes precedence over the `metric-out` statement configured under the neighbor configuration for the same group. [PR/31848]
- The CLI allows you to commit a configuration that specifies a value higher than 32 for the `metric` statement at the `[edit protocols dvmrp interface all]` hierarchy level, but values higher than 32 are invalid. [PR/33429]

- If you configure the `sham-link` statement at the `[edit routing-instances instance-name protocols ospf area]` or `[edit routing-instances instance-name protocols ospf]` hierarchy level on a provider edge (PE) router, extraneous OSPF link-state advertisements (LSAs) might be added. In some cases, this can result in a routing loop between the customer edge (CE) and PE routers. [PR/40000]
- When you configure damping globally and use the import policy to prevent damping for specific routes, and a new route is received from a peer with the local interface address as the next hop, the route is added to the routing table with default damping parameters, even though the import policy has a nondefault setting. As a result, damping settings do not change appropriately when the route attributes change. [PR/51975]
- When you issue the `show ldp traffic-statistics` command, the following system log message might be generated for all forwarding equivalence classes (FECs) with an ingress counter set to zero: “send rnhstats GET: error: ENOENT -- Item not found.” [PR/67647]
- When routes are propagated across IBGP, the `show bgp group statistics` command output does not display AS numbers correctly. [PR/69098]
- If ICMP tunneling is enabled on the router and you configure a new logical router that does not have ICMP tunneling enabled, the feature is globally disabled. [PR/81884]
- If ICMP tunneling is enabled on the router and you configure a new logical router that does not have ICMP tunneling enabled, the feature is globally disabled. [PR/81884]
- When you specify a link-local interface for the `interface` statement at the `[edit routing-options rib inet6.0 static route address/mask-length qualified-next-hop address]` hierarchy level, the commit operation fails with the message `RT: next-hop interface-name is not point-to-point`. [PR/99293]
- When the flow of multicast traffic changes because an OSPFv3 link goes down, the output from the `show multicast statistics inet6` command reports incorrect values in the `In kbytes` and `In packets` fields for the new ingress interface. [PR/234969]
- Access-Internal routes are not entered in to the forwarding table for unnumbered Ethernet interfaces. [PR/252220]
- The address for the flow route is terminated at 348 characters. It is a cosmetic issue and affects the flow route display in `show route`. [PR/273385]
- Multicast Source Discovery Protocol (MSDP) incorrectly reports a non-existent security association (SA), resulting in the SA remaining in Protocol Independent Multicast (PIM) when it is deleted in the MSDP. [PR/277310]

## MPLS Applications

- If you configure a label-switched path (LSP) with the `no-cspf` statement at the `[edit protocols mpls]` hierarchy level, the LSP might cycle up and down several times before stabilizing. [PR/10415]
- If a circuit cross-connect (CCC) traverses a forwarding adjacency (FA) label-switched path (LSP), traffic forwarding might be affected. [PR/60088]

- RSVP graceful restart does not function for LSPs that have a forwarding adjacency (FA) label-switched path (LSP) as a next hop. [PR/60256]
- When you enable per-packet load balancing on parallel label-switched paths (LSPs), the output of the `show mpls lsp ingress` command might display all the routes on only one of the LSPs even when traffic is evenly balanced across the LSPs. [PR/70487]
- On M-series and T-series routing platforms, if MPLS traffic is being forwarded on the secondary path of an LSP when the primary path is also functional, the **Traffic statistics** section of the output from the `monitor label-switched-path lsp-name` command might show incorrect values. [PR/80591]
- The `show mpls lsp detail` command does not display an LSP's setup and hold priorities (the `Priorities` field is omitted) if they are set to their default values, even if the defaults are set explicitly at the `[edit protocols mpls label-switched-path path-name priority]` hierarchy level. As a workaround, issue the `show mpls lsp defaults` command to display the priority values. [PR/103128]
- On an M120 router, the `ping mpls rsvp` command fails when an LSP is configured for link protection (the `link-protection` statement is included at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level) and traffic is being routed through the bypass LSP. [PR/233693]
- In the output from the `show mpls lsp` command, the column labeled **ActivePath** is about 16 characters wide. When the name of an LSP path is longer than that, subsequent values on the line do not align correctly with their headers. [PR/237229]
- When more than 5 link-protected or node-link-protected label-switched paths (LSPs) to the same destination are used with per-packet load balancing, some bypass next-hops might not be included in the active route. This can occur after a primary link flap. [PR/259219]
- On M-series and T-series routers, when an MPLS LSP is optimized, the MPLS MIB counters associated with the path change event are not updated. [PR/265931]
- Sometimes a traffic engineered label-switched path (LSP) remains up when it should go down. [PR/300919]
- When a Layer 2 circuit comes back up after a disruption, it remains attached to the old label, so traffic does not pass through the Layer 2 circuit connection. [PR/306043]

## VPNs

- When you modify the `frame-relay-tcc` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level of a Layer 2 VPN, the connection for the second logical interface might not come up. As a workaround, restart the chassis process (`chassisd`) or reboot the router. [PR/32763]
- When VPLS nonstop active routing is enabled and you modify the VPLS instance (for example, change the instance type or its route distinguisher), the routing process (`rpd`) might stop and the system might produce a core dump. [PR/231234]
- Traffic might not flow when an ATM interface is used as the access circuit on an M120 router. [PR/255160]

## Class of Service

- When you configure an ES PIC, a message similar to the following might be written to the system log: “fpc0 LCHIP(3): Unable to fathom what channel used by IFD *id*.” There is no operational impact. [PR/36184]
- If you deactivate or activate an aggregated Ethernet interface, the Packet Forwarding Engine might report errors. [PR/50090]
- When a logical tunnel (lt) interface is the outbound interface, JUNOS software does not support the IEEE 802.1p rewrite rule. [PR/55903]
- If you try to configure a scheduler map containing two forwarding classes that are mapped to the same queue, the class-of-service scheduler is not applied to the Packet Forwarding Engine. As a workaround, configure a single forwarding class for each available queue. [PR/57907]
- On M-series routers connected by VLAN circuit cross-connects (CCCs) and configured with class of service (CoS), when explicit forwarding (EF) traffic is generated from the ingress customer edge router (CE1) to the egress customer edge router (CE2), the ingress provider edge router (PE1) properly marks the packets with default EXP bits and sends the packets out queue 1, but the intermediary core router forwards all traffic through queue 0 instead of sending it through the EF queue. As a workaround, include the `no-control-word` statement at any of the following hierarchy levels: [edit logical-routers *logical-router-name* protocols l2circuit neighbor *address* interface *interface-name*], [edit protocols l2circuit neighbor *address* interface *interface-name*], [edit logical-routers *logical-router-name* routing-instances *routing-instance-name* protocols l2vpn], or [edit routing-instances *routing-instance-name* protocols l2vpn]. [PR/65280]
- When you configure a specific classifier for a logical unit, it does not override the fixed classifier configured using wildcards. [PR/68888]
- Adding and deleting an interface many times by configuring the `scheduler-map-chassis` statement at the [edit class-of-service interfaces *ge-1/1/port*] hierarchy level might cause a memory leak in the class-of-service process (`cosd`). As a workaround, restart the class-of-service (CoS) process (`cosd`). [PR/82546]
- If you configure CoS traffic control profiles on every logical interface by using the `*` wildcard to represent the interfaces, the configuration cannot be committed. In other words, the commit operation fails if you include the `input-traffic-control-profile` and `output-traffic-control-profile` statements at the [edit class-of-services interfaces *type-fpc/pic/port \**] hierarchy level. [PR/100690]
- On MX-series routers, when you configure VPLS over an LSI interface, classification does not work on the egress PE router for traffic flowing from the core of the network to the egress CE router. [PR/240777]
- If you configure the `tri-color` statement at the [edit class-of-service] hierarchy level, the drop counters for the `show interfaces queue` command appear to not work for the medium-high (yellow) priority traffic and the low (green) priority traffic. The drop counter for the high priority traffic (red) functions normally. [PR/258499]
- In JUNOS Release 8.4 and later, the 'commit' or 'commit-check' operation fails if a rewrite rule is defined both at the [edit class-of-service interfaces *interface-name*unit *logical-unit-number* rewrite-rules] hierarchy level and in a configuration group (defined at the [edit groups] hierarchy level) that is applied

to that interface. The correct behavior is for the directly applied rule to override the rule inherited from the configuration group. [PR/261229]

- The output from the `show class-of-service interface interface-name` command includes the `Input scheduler map` field even when you configure egress-only mode for the PIC that houses the interface (by including the `mode egress-only` statement at the `[edit chassis fpc slot-number pic slot-number traffic-manager]` hierarchy level). [PR/275038]

### Forwarding and Sampling

- On M320 and T-series routing platforms, when you configure interface output sampling, packets sometimes might travel through the output firewall. As a workaround, configure a firewall filter on the output interface with the `then sample` statement and the `then next term` statements. The workaround provides the same functionality as the other configuration, but avoids the problem behavior. [PR/70473]

### Routing Policy and Firewall Filters

- The extended Dynamic Host Configuration Protocol (DHCP) relay agent feature does not function properly on a nondefault logical router. This means that although the JUNOS CLI permits you to include the `dhcp-relay` statement at the following hierarchy levels, the feature does not work properly when you do so:
  - `[edit logical-routers logical-router-name forwarding-options]`
  - `[edit logical-routers logical-router-name routing-instances]`
  - `[edit logical-routers logical-router-name routing-instances routing-instance-name forwarding-options]`

[PR/82275]

- On MX-series routers running JUNOS Release 8.4 and later, entries in the MAC address table expire three times faster than on MX-series routers running JUNOS Release 8.3 and earlier, and on M-series and T-series routing platforms running any release of the JUNOS software (including JUNOS Release 8.4 and later). To configure the correct effective value on MX-series routers running JUNOS Release 8.4 and later, specify a value for the `mac-table-aging-time` statement at the `[edit protocols l2-learning]` hierarchy level that is three times the desired value. For example, if you want the expiration time to be 15 seconds, specify 45 seconds. [PR/241485]

### Network Management

- The following groups of MIB objects do not segregate the data they return according to the routing instance specified in an SNMP request: `vrmpMIB`, `jnxCosIfqStatsTable`, and `jnxCosQstatTable`. [PR/63045]
- When you commit a configuration that includes the `max-queues-per-interface` statement at the `[edit chassis fpc slot pic slot]` hierarchy level, the MIB II process (`mib2d`) might generate a core file and stop operating. [PR/99197]

- If an element number in an MIB object's OID is greater than 2,147,483,647 (2 to the 31st power, minus 1), the `snmp mib walk` and `snmp mib get` commands fail. [PR/237856]

## Previous Releases

---

### 8.5R3

The following issues have been resolved since JUNOS Release 8.5R3. The identifier following the description is the tracking number in our bug database.

#### Software Installation

- When you issue the `request system partition hard-disk` command, the hard disk repartition fails and the disk becomes unusable. The disk can be recovered by taking a snapshot from the compact flash card and rebooting the router. [PR/269493: This issue has been resolved.]

#### Platform and Infrastructure

- If TCP and UDP probe servers (configured at the `[edit services rpm]` hierarchy level) are configured on a router and a TCP and a UDP probe is made to it simultaneously from some other router, multiple occurrences of the `RMOPD_SENDMSG_FAILURE` message are generated and recorded in the system log. [PR/66570: This issue has been resolved.]
- During a Routing Engine switch over, the Flexible PIC Concentrator (FPC) might reset multiple times. [PR/70857: This issue has been resolved.]
- ARP records learned in a VPN routing and forwarding (VRF) instance are not cleared when the peer interface goes down. [PR/82247: This issue has been resolved.]
- If a small form-factor pluggable transceiver (SFP) does not respond to a request for diagnostic data, a message is written to the system log. The message is unnecessary because the failure to respond has no operational impact. [PR/97718: This issue has been resolved.]
- On T-series platforms, if you include the `indirect-next-hop` statement at the `[edit routing-options forwarding-table]` hierarchy level for VPN routes, routing ASIC SRAM utilization increases by approximately 30 percent or 8 bytes per route. [PR/98738: This issue has been resolved.]
- Due to changes in the JUNOS TCP/IP networking stack, the output of the `show connections` CLI command may be different from JUNOS 8.4 and earlier. [PR/103330: This issue has been resolved.]
- PFE might incorrectly log an error (described in the PR). This might happen when there is some sort of aggregate member link status change. [PR/105841: This issue has been resolved.]
- When IPSec is configured on a logical interface and the protocol family is IPv6, graceful Routing Engine switchover (GRES) might fail if an MTU change is attempted on that interface. [PR/230128: This issue has been resolved.]

- J-Series multilink interfaces operate correctly when fragments are in round-robin fashion and arrive in order. However, if the fragments are out of order, then they suffer some latency and packet-loss during reassembly. [PR/240019: This issue has been resolved.]
- When the management interface (fxp0) initialization does not complete, the interface loses network connectivity and does not respond to any commands within the timeout period of 10 milliseconds. [PR/253479: This issue has been resolved.]
- On some Routing Engines, the smartd process may display the following error: “atastandbyarmset”. [PR/253775: This issue has been resolved.]
- Using the “-ox” option with the smartd process is not recommended on mounted devices because it may result in unexpected behavior. [PR/255473: This issue has been resolved.]
- When an address rename operation is performed on Gigabit Ethernet interfaces, filters are removed and then added back. The operation can sometimes be replicated to the backup Routing Engine as a single change. In this scenario, the backup Routing Engine attempts to delete the filter and add it back using the index specified by the master Routing Engine. However, the entry is not deleted, leading to a mismatch in the index usage between the master and slave Routing Engines, which causes the Routing Engine to produce a core file and stop operating. [PR/258927: This issue has been resolved.]
- The UDP ping server does not respond to probes sent through a routing instance other than the default (inet.0). [PR/260097: This issue has been resolved.]
- A multicast enabled router, configured with aggregated Ethernet, or aggregated SONET, interfaces could experience a kernel crash, when the constituent, or aggregate interfaces go down. [PR/264579: This issue has been resolved.]
- When you change interface configuration from point-to-point encapsulation to Frame Relay encapsulation, the routing platform kernel might generate a core file and stop operating. [PR/265025: This issue has been resolved.]
- When IPv6 traffic is present on MLPPP bundles using MultiServices PIC, the Service PIC may core due to excessive logging as a result of invalid checksum calculation. [PR/266214: This issue has been resolved.]
- If there are many aggregate next hops and BGP routes pointing at some of them, a quick link flap combined with the BGP route churn might cause the Packet Forwarding Engine to restart unexpectedly. [PR/268204: This issue has been resolved.]
- With certain traffic patterns, MX-series and M320 routers with 3.0 forwarding ASICs might experience packet loss. To recover, you must reboot the affected DPC. [PR/268274: This issue has been resolved.]
- When FPC is brought online immediately after being taken offline, it may not be properly initialized due to a race condition in timing between the FPC and Routing Engine. [PR/272086: This issue has been resolved.]
- In JUNOS software Release 8.5R1 only, if the router receives an MPLS LSP ping packet, it may cause a kernel memory leak in the network packet buffer and upon exhaustion no packet transfer is possible between the Routing Engine and the Packet Forwarding Engine (PFE). There is no workaround. [PR/273024: This issue has been resolved.]

- Using the set forwarding-options sampling output file filename `sampled.pkts`, the `sampled.pkts` file (which should be created in `/var/tmp` in the router) is not being created. If restart sampling is called, the file gets created. This is because the j-flow license database is not working properly when sampling is configured. A workaround for this issue is to issue the `show system license usage` command. The license database is then updated and the file is created. [PR/275473: This issue has been resolved.]
- When both BPG multipath and indirect-next-hop are turned on, the next hop referenced by the prefixes for BGP multipath will not be installed in the routing chip, resulting in packet drop. [PR/275586: This issue has been resolved.]
- A large virtual memory size (> 700 MB) is reported for daemons that access the JUNOS configuration database in the output of the `show system processes extensive` command. The reason for the increase in virtual memory size is because 700 MB of the virtual space is pre-reserved for the JUNOS configuration database. This reserved space is just a placeholder in the daemon's virtual address space. Customers do not need to be concerned by this increase in virtual memory size because the actual amount of physical memory consumed by the configuration database is not affected. [PR/276378: This issue has been resolved.]
- In scenarios where the PFEs are restarted quite frequently, the PFE might crash. Two causes have been identified. One is a race condition with the PIC state, and the other has to do with duplicate family add messages. [PR/276539: This issue has been resolved.]
- When packets are queued for several seconds due to interface congestion, in some cases, packet CRC errors are reported. In other cases, the egress interface stops forwarding traffic (either all traffic is halted or only packets larger than 320 bytes are not forwarded). As a workaround, configure queues with a transmit rate of at least one percent of the line rate. When using strict, high-priority queues, include policers to prevent interface congestion and the starving of lower-priority queues. Another alternative is to use priority, high instead of strict, high queues. [PR/277853: This issue has been resolved.]
- Once protocol tracing is enabled, IO writes to the hard drives might get blocked and daemons sending packets will get behind. This applies only to JUNOS software Release 8.5R2 or higher. [PR/278580: This issue has been resolved.]

### User Interface and Configuration

- TACACS+ accounting start/stop requests are incompatible with Cisco ACS. The fix is to add the `no-cmd-attribute-value` statement at the `[edit system tacplus-options]` hierarchy level. When this is enabled, the JUNOS software sets the value of the `cmd` attribute in TACACS+ accounting start and stop requests to a null string. This is the behavior Cisco ACS expects in order to save accounting requests to the Accounting file; otherwise, the requests are saved to the Administration file. [PR/252472: This issue has been resolved.]
- Changing any unrelated configuration causes the `apply-groups` statement under routing instances to trigger an RPD reinitialization during commit even when the change didn't match any attributes in the group. Also, `rollback` performs the equivalent of `load override` instead of `load update`. This causes the routing protocol process to be reinitialized during commit. [PR/259740: This issue has been resolved.]

- On upgrading to 8.5R1.13, the MOTD might be displayed twice. [PR/268625: This issue has been resolved.]
- If you issue the `copy` command in edit mode for configurations under the `[edit groups]` hierarchy, the `mgd` process exits and dump core. [PR/269034: This issue has been resolved.]
- Commands such as `rollback <#>` or `show | compare rollback <#>` take several minutes to complete if the configuration has a policy statement with a term that contains a large number of route-filter entries (in the order of thousands). [PR/272350: This issue has been resolved.]
- The `system syslog source-address` change does not take effect. As a workaround, you can restart the `eventd` process (or reboot the Routing Engine). [PR/272434: This issue has been resolved.]
- JUNOS devices cannot be managed by Session and Resource Control (SRC) software. [PR/273117: This issue has been resolved.]
- When you use the `configure private` command and then change two static routes at the same time at the `[routing-options static]` hierarchy level, the commit may fail. As a workaround, configure one static route at a time. [PR/273251: This issue has been resolved.]
- Issuing the `replace pattern` command may terminate the telnet session with `mgd` core due to a function not being invoked correctly. [PR/274830: This issue has been resolved.]
- A user without maintenance permission cannot `su` to a non-root user from the shell. [PR/277888: This issue has been resolved.]
- Login class permissions might not work properly for JUNOS software Releases 8.4R1 and higher. [PR/278950: This issue has been resolved.]

## Interfaces and Chassis

- On M20 routers, when you start the router with Routing Engine 0 and System and Switch Board (SSB) 0 as master components, issue the `request chassis routing-engine master switch` command, and then log in to Routing Engine 1 and issue the `request chassis ssb master switch` and `request system reboot` commands, the “ONLINE” LED might remain lit on both SSBs. [PR/74283: This issue has been resolved.]
- When you configure point-to-multipoint Frame Relay, the router might generate a core file. [PR/82303: This issue has been resolved.]
- The `show vrrp` commands which were earlier available only from the top level, has been made available from an LR context as well. [PR/253956: This issue has been resolved.]
- When you issue the `request chassis routing-engine master switch` command to change Routing Engine mastership, the `jnxRedundancySwitchover` SNMP trap is not generated. (However, the event is recorded in the system log or chassis process log file if logging is appropriately configured.) [PR/254637: This issue has been resolved.]
- Configuring an interface IPv6 address as preferred or primary will generate a log message similar to the following: "DCD\_CONFIG\_WRITE\_FAILED: Interface

'interface-name' configuration write failed for an IFA CHANGE: Operation not supported". [PR/258531: This issue has been resolved.]

- When you have IPv6 enabled, the IPv6 input transit counter might not work correctly. [PR/260704: This issue has been resolved.]
- Small correction in printing the "expected" value in VRRPD\_VIP\_COUNT\_MISMATCH message for an IPv6 address. [PR/261415: This issue has been resolved.]
- When the committed DLCIs are in the process of coming up, if there are any configuration changes on that specific interface, the DLCIs that are down (or have not yet come up) at that time, will not come up until that IFD is reset (offline and online) or until the logical interfaces are deactivated and activated. If any configuration change is made to the Frame Relay interface when the existing DLCIs that are down, they will not come up until that IFD is reset (offline and online) or until the logical interfaces are deactivated and activated. [PR/261501: This issue has been resolved.]
- When telnet/ssh sessions are negotiated on the router, if enough packets are lost to cause TCP SACK (Selective ACK) to be exercised, SACK errors appear in the logs and might cause the FPCs to disconnect from the Routing Engine. [PR/265957: This issue has been resolved.]
- Input statistics for aggregated Ethernet interfaces incorrectly report zero, regardless of input traffic volume. [PR/266271: This issue has been resolved.]
- The chassis LED status returned by the MIB jnxLEDState does not reflect the actual chassis alarm LED. [PR/266326: This issue has been resolved.]
- SFPCs caused the warning message "WARNING: Unknown FPC 0x1f4" when checking PIC compatibility. [PR/266854: This issue has been resolved.]
- When there is an Address Resolution Protocol (ARP) entry for Virtual IP (VIP), the Virtual Router Redundancy Protocol (VRRP) might not respond to ARP requests for VIP while transitioning to master state. [PR/268627: This issue has been resolved.]
- An error in the chassis process (chassisd) caused a small memory leak when the **show chassis hardware extensive** command was executed. [PR/268925: This issue has been resolved.]
- On an MX platform, when there is a lot of traffic going to the Routing Engine and the route changes, the Dense Port Concentrator (DPC) might trigger an assertion without dumping core. [PR/269699: This issue has been resolved.]
- When Routing Engine mastership is repeatedly switched, routing information maintained on the master and backup Routing Engine might be out of sync, causing all Packet Forwarding Engines to reset. [PR/271141: This issue has been resolved.]
- When you configure multiple interfaces with **vlan-id-range** and cover a large number of VLAN IDs, the DPC might restart unexpectedly. [PR/271456: This issue has been resolved.]
- Clocking configuration under **so-x/y/z** level may conflict with clocking under **coc-x/y/z** level. Customer should only configure clocking under **coc-x/y/z**. [PR/272920: This issue has been resolved.]

- For RLSQ multilink feature with GRES, any addition deletion of the member T1/E1 links after GRES switchover will not operate correctly. The work around is to deactivate/activate the entire RLSQ bundle when links are deleted/added after a GRES switchover. [PR/273528: This issue has been resolved.]
- When a shape-rate is configured to a logical interface on IQ2 Ethernet PICs, it may not be applied, and traffic exceeding the preconfigured rate may still go through. [PR/273831: This issue has been resolved.]
- On Ethernet IQ2 PICs, if you change the configuration at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level from forwarding-class to classifiers, the behavior aggregate (BA) classification table mapping might be incorrect. As a workaround, deactivate the virtual circuit and then reactivate it. [PR/275693: This issue has been resolved.]
- On M320 Enhanced 3 FPC, when next term is used in firewall filter term, the firewall filter may block all traffic on interface where the filter is applied. [PR/278325: This issue has been resolved.]
- In VPLS instance on MX series routers, if an interface is configured as nontagged, vlan-maps will not process frames with dual tags. [PR/279669: This issue has been resolved.]

### Services Applications

- If Network Address Port Translation (NAPT) is configured and multiple short-lived flows are established, ports on AS PICs might not be assigned correctly. In some cases, this situation causes the AS PIC to stop functioning. [PR/95019: This issue has been resolved.]
- Flow monitoring version 9 MPLS and MPLS-IPv4 templates do not work on the TX matrix routing platform or on dual Routing Engine routing platforms configured with graceful switchover. When you configure flow monitoring version 9 on a TX matrix routing platform, LCCs will crash. When you configure flow monitoring version 9 on a router configured for graceful switchover, issuing the commit synchronize command will cause the backup Routing Engine to crash. [PR/98372: This issue has been resolved.]
- If Network Address Port Translation (NAPT) is configured and multiple short-lived flows are established, ports on AS PICs might not be assigned correctly. In some cases, this situation causes the AS PIC to stop functioning. [PR/229287: This issue has been resolved.]

## General Routing

- Community expression can be mistakenly interpreted as an extended community instead of a regular community regex. [PR/251510: This issue has been resolved.]

## Routing Protocols

- When determining which BGP route to prefer, the JUNOS software does not compare route characteristics in the order specified by RFC 4456, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)." [PR/70687: This issue has been resolved.]
- When two redundant PE routers send multicast traffic towards their MVPN backbone, the JUNOS software might not perform the actions necessary to prevent traffic duplication. [PR/72447: This issue has been resolved.]
- The RPD process may core after upgrading to 7.6R2.6 when upgraded from 7.3R2.9. There is no workaround as of now to avoid the problem. [PR/77495: This issue has been resolved.]
- When a new receiver joins the tree between the source and RP for which there is already a receiver for the same group upstream to the RP, there might be delay for the traffic to be received by the new receiver until the next Register message is received. [PR/228708: This issue has been resolved.]
- When a new PIM neighbor is discovered on an IPv6 network, the appropriate system log message was not filed. [PR/230342: This issue has been resolved.]
- When PIM receives an (\*,G) leave and an (S,G) join for an (S,G) entry that has been pruned and when the neighbor entry is deleted, the JP state for the (S,G) join is not deleted, resulting in PIM going in to an infinite loop. [PR/235978: This issue has been resolved.]
- The `show bgp summary` command output changed to show BGP IPv6 neighbors in sorted order, based on the IPv6 address. [PR/237127: This issue has been resolved.]
- In rare cases, the routing protocol process (rpd) may have an iflist that is different from the kernel's. A subsequent route add may try to use the nhindex of this iflist. The kernel returns an error and the route sits in krt's queue forever. [PR/252489: This issue has been resolved.]
- When you deactivate a large number of routing instances (about 335 or more), LDP sessions might go down and the following message might be written to the system log: "RPD\_PPM\_WRITE\_ERROR: ppm\_send: write error on pipe to ppm (Broken pipe)." [PR/260477: This issue has been resolved.]
- If you issue the `show route extensive` command while the routing protocol process (rpd) is resolving BGP routes, the process might generate a core file and restart. [PR/260527: This issue has been resolved.]
- ARP entries created by DHCP relay are not programmed as permanent. Therefore, any other ARP request from different users on the same Ethernet segment can overwrite this entry at any time. [PR/264332: This issue has been resolved.]
- For VPLS, when the router-distinguisher is changed on an existing instance, the RPD may be running at high CPU utilization and not responding to the CLI. [PR/270204: This issue has been resolved.]

- The redistribution of OSPF point-to-point (P2P) LAN interfaces from within a routing instance in to the main routing instance can fail, displaying the following log message: “Jan 23 15:24:50 router rpd[6063]: cannot perform nh operation ADDANDGET nhop 0.0.0.0 type unicast index 0 errno 45.” As a workaround, redistribute the interfaces routes within the routing instance in to the main routing instance. [PR/271130: This issue has been resolved.]
- The routing process can restart unexpectedly if it receives a BGP flow NRLI specification with an undefined subcomponent type. [PR/274421: This issue has been resolved.]
- MSDP is reporting that an SA is active when the entry no longer exists. This is why the entry lingers in PIM once it has been deleted in MSDP. [PR/277310: This issue has been resolved.]
- AS path recorded information was being shown at the wrong place. [PR/281023: This issue has been resolved.]
- With a minimum of two PIM ASM routers and receivers on a LAN segment, if interface MISMATCH errors occur, then an s,g entry will not be created. This is applicable for both native PIM configurations as well as multicast VPN configurations. [PR/281662: This issue has been resolved.]

### MPLS Applications

- On M series and T series routers, the output of the `monitor label-switched-path lsp-name` command might show incorrect LSP statistics when the both the secondary and primary path of an LSP are UP and traffic is being forwarded on the secondary path. [PR/80591: This issue has been resolved.]
- The `jinstall ugrade` error message “task\_get\_port: getservbyname("ldp", "tcp") failed, using port 646” is a minor warning which can be ignored. However a fix has been checked in for this issue. [PR/102209: This issue has been resolved.]
- A router with a point-to-multipoint transmit-switch connection configured might stop functioning if the transmit label-switched path of the connection flaps. [PR/229175: This issue has been resolved.]
- MVPN P2MP deactivating of the VT interface in a particular VPN, say VPN-A on receiver PE, affects multicast traffic forwarding in other VPNs for a few seconds. Same issue is seen during the activate sequence of the VT interface in VPN-A. [PR/252697: This issue has been resolved.]
- The order of addresses returned as the ERO for an LSP by the SNMP `mplsPathExplicitRoute` object is reversed from the correct order displayed as the `Computed ERO` in the output of the `show mpls lsp extensive` command. [PR/263462: This issue has been resolved.]
- When graceful restart is enabled (by the `graceful-restart` statement at the `[edit routing-options]` hierarchy level) and an interface is configured as an egress interface in a CCC switch for a point-to-multipoint LSP (the `output-interface` statement is included at the `[edit protocols connection p2mp-receive-switch]` hierarchy level), if the LSP flaps repeatedly, the interface might stop forwarding traffic. [PR/264930: This issue has been resolved.]

- The system can leak next-hop resources when RSVP link protection is enabled. [PR/265295: This issue has been resolved.]
- On M series and T series routers, when an MPLS LSP gets optimized, the MPLS MIB counters associated with the path change event are not updated. [PR/265931: This issue has been resolved.]

## VPNs

- When configuring multicast VPN for point to multipoint, the `tunnel-limit` statement for dynamic selective provider tunnels is not functioning. [PR/250701: This issue has been resolved.]
- A receiver directly attached over an Ethernet connection to a sender PE configured with multicast VPN over point to multipoint fails to receive multicast traffic. No such issue is observed when the receiver is connected over a Sonet or ATM link. [PR/252314: This issue has been resolved.]
- The VT tunnel interface is not recreated upon reboot of Tunnel PIC or FPC reboot holding the Tunnel PIC. As a workaround, deactivate and activate the VT tunnel interface stanza in the interface configuration. [PR/266170: This issue has been resolved.]

## Class of Service

- When you configure the EXP classifiers to a routing instance, the class-of-service (CoS) process (`cosd`) might dump core. [PR/101490: This issue has been resolved.]
- Configuring `huge-buffer-temporal` in CoS may cause commit to fail when using `configure private` [PR/265762: This issue has been resolved.]
- If a rewrite rule is not defined at the `[edit class-of-service rewrite-rules]` hierarchy level for every forwarding class defined at the `[edit class-of-service forwarding-classes]` hierarchy level, when the class-of-service process restarts (for example, when there is a graceful Routing Engine switchover), the process does not initialize internal data structures correctly. As a workaround, define a rewrite rule for every forwarding class. [PR/268541: This issue has been resolved.]
- On MX960 routers, the class-of-service process does not provide information about SNMP objects whose names begin with "jnxCosQstat." As a result, SNMP queries on those objects fail with an error message. [PR/269419: This issue has been resolved.]
- On Type 4 FPC or T1600 platforms if you change the exp rewrite rules it may trigger an overflow of the L2 programs which will result in incorrect programmed next hop entries and traffic is not forwarded. You need reboot the FPC to recover. [PR/279625: This issue has been resolved.]

## Forwarding and Sampling

- On M120 and MX-series routers, if you configure both a firewall filter and interface sampling for ingress traffic on the same interface (by including both the `filter` and `sampling` statements at the `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy level), the interface discards all incoming packets. As a workaround, implement input sampling as an action in the `then` section of a firewall filter. [PR/103206: This issue has been resolved.]
- The implicit DHCP firewall filter is removed from all interfaces once other configuration changes are performed at the firewall or interface level. [PR/261009: This issue has been resolved.]
- Firewall filters that include the `source-address` statement at the `[edit firewall filter filter-name term term-name from]` hierarchy level might not process traffic correctly. As a workaround, reorder the terms in the filter. [PR/262491: This issue has been resolved.]
- If you deactivate the `nonstop-bridging` statement at the `[edit protocols layer2-control]` hierarchy level and then a graceful Routing Engine mastership switchover occurs, the bridge ID for a routing instance might be reset to all zeros and the port state for some component interfaces in the routing instance might change to "Blocking." Both conditions are reported by the `show spanning-tree bridge routing-instance detail` command. [PR/264982: This issue has been resolved.]
- If a term in firewall filter specifies a range of values for source or destination address or port, the filter might not match packets as expected. As a workaround, define the addresses and ports explicitly. [PR/265023: This issue has been resolved.]
- When a firewall policer with loss-priority action is used under multiple interfaces, only the first instance of the interface policer will have the right PLP setting behavior. [PR/274346: This issue has been resolved.]
- At the `[edit routing-options flow route route-name]` hierarchy level, when a range of numeric values is selected for firewall filtering using a statement, the match condition is ignored. Such statements typically include `destination-port`, `port`, `protocol`, and `source-port`. [PR/275650: This issue has been resolved.]

## 8.5R2

The following issues have been resolved since JUNOS Release 8.5R2. The identifier following the description is the tracking number in our bug database.

### Platform and Infrastructure

- The MultiServices PIC might not work correctly when the PIC is loaded and frequent commands related to the PIC are issued. [PR/81826: This issue has been resolved.]
- When IPsec is configured on a logical interface and the protocol family is IPv6, graceful Routing Engine switchover (GRES) might fail if an MTU change is attempted on that interface. [PR/230128: This issue has been resolved.]
- J-series multilink interfaces behave well when fragments are in round-robin fashion and arrive in order. However, if fragments are out of order, then they

will suffer some latency and packet loss during reassembly. [PR/240019: This issue has been resolved.]

- When using file copy FTP, the IP address specified is the source address is not used for establishing a connection with the peer FTP server. [PR/240580: This issue has been resolved.]
- When graceful switchover and RLSQ interfaces are configured, the kernel generates a core file on the backup Routing Engine with the error message "panicstr: rnh\_index\_alloc: nhindex 116435 could not be allocated." [PR/241502: This issue has been resolved.]
- When an unnumbered Ethernet interface has a loopback address as a donor and if the address configured on the loopback is a subnet address, a ping to the subnet address does not work. [PR/253804: This issue has been resolved.]
- In some situations, the interface counter account doubles the number of packets. [PR/253946: This issue has been resolved.]
- The `clear arp` command does not function for logical routers. [PR/253957: This issue has been resolved.]
- The `show arp` command does not function for logical routers. [PR/253958: This issue has been resolved.]
- A router running multicast over aggregate SONET or aggregate Ethernet interfaces could experience a Packet Forwarding Engine crash when a constituent link flaps. [PR/257691: This issue has been resolved.]

### User Interface and Configuration

- If you use telnet to connect to the JUNOScript Perl module, the connection fails if the password or login name includes special characters. [PR/241236: This issue has been resolved.]
- Certain JUNOScope wizards (devices, groups, users, schedules, and RADIUS configuration ) do not work with Netscape 7.0. The workaround is to use Netscape 6.2 [PR/260326: This issue has been resolved.]

### Interfaces and Chassis

- If a low-speed bundle is congested, the jitter for link fragmentation and interleaving (LFI) traffic is high even though fragmentation is configured for Multilink Point-to-Point Protocol traffic. [PR/77862: This issue has been resolved.]
- If you clear IPv6 statistics, deactivate IPv6 route accounting at the `[edit forwarding-options family inet6 route-accounting]` hierarchy level, and resume IPv6 traffic across an interface, the "Input bytes" and "Input packets" fields in the output of the `show interfaces extensive` command might display incorrect values. [PR/99461: This issue has been resolved.]
- On the M120 and MX-series platforms and the M320 Enhanced III FPCs only, forwarding IPv6 transit packets might stop transmitting traffic but still be able to receive traffic. All packets going out this interface will be dropped. To recover, you need to reboot the FPC on M320 platforms, the DPC on the MX-series, or the FEB on M120 platform. There is no workaround. [PR/105266: This issue has been resolved.]

- For Gigabit Ethernet interfaces on J-series Services Routers, the `link-mode` and `speed` statements at the `[edit interfaces ge-fpc/pic/port]` hierarchy level are mutually dependent; that is, if you include one, you must include the other. If you do not, the interface process generates a warning and uses autonegotiated values. For Gigabit Ethernet interfaces on other routing platform types, the `speed` statement is not available, so including the `link-mode` statement alone is valid. Nevertheless, the interface process writes the following message to its log and the system log: "Speed and linkmode duplex settings are mutually required." (Note further that the `ink-mode` statement is actually nonoperational on non-J-series routing platforms, because the only valid value for it is the default, 'full-duplex'.) [PR/228857: This issue has been resolved.]
- The "accept data" warning message for VRRP might not correctly display the logical unit identifier. [PR/236135: This issue has been resolved.]
- When a PIC detachment process takes a long time for an IQ2 PIC (for example, when a large number of route updates are triggered after an IQ2 PIC is brought offline), the PIC chassis process (`pic-chassisd`) connection might not be closed properly. [PR/239944: This issue has been resolved.]
- On the J4350 and J6350 Services Routers, the `show chassis interfaces extensive` command does not display "Carrier Transitions" and certain other statistics properly. [PR/241086: This issue has been resolved.]
- Padded MPLS-encapsulated IPv4 packets that exit an LSP can cause the egress interface to stop forwarding packets. This can happen when the router is configured as a VPN PE, or when the router is the penultimate node of an LSP. The problem only occurs when the packet has been padded to meet the minimum Layer 2 frame size (for example, Ethernet media require frames to be a minimum of 64 bytes long). This issue applies to the M120 and MX-series platforms and to the M320 Enhanced III FPCs. To recover, reboot the FPC on the M320, the DPC on the MX-series, or the FEB on M120 routing platforms. [PR/251042: This issue has been resolved.]
- The ATM PIC driver might not always use the minimum port shaping rate (of all the ports on a multiport ATM DS3 or E3 PIC) selected for cell transmission shaping, in situations where the DS3 or E3 port parameters are not identical on all ports of a multiport ATM DS3 or E3 PIC. The PIC shaping rate is always updated to conform to the last port setting updated by the PIC software driver, rather than using the minimum port (shaping) rate. There is no syslog message to inform the user of the shaping rate decision applied by the software driver. [PR/252837: This issue has been resolved.]
- On the J-series Services Routers, on 4-port Fast Ethernet Enhanced Physical Interface Modules (ePIMs), interfaces might stop working correctly and lock up a port when operating in half-duplex mode. As a workaround, hard code the `link-speed` and `link-mode` to 100m full duplex. [PR/253329: This issue has been resolved.]
- When you configure a logical interface with encapsulation `ether-vpls-over-atm-llc`, the packets destined for to the Routing Engine are dropped by the Packet Forwarding Engine. [PR/255713: This issue has been resolved.]
- When a Routing Engine assumes mastership, it attempts to reconnect to the Packet Forwarding Engine. The Packet Forwarding Engine sends all the information to the master Routing Engine. The new master Routing Engine then

attempts to retrieve SFP information from the PIC, but the PIC fails to send it. [PR/256032: This issue has been resolved.]

- On a M40e router with 4 MS00 PICs and 4 ChOC12, the SONET interface might display remote defect indication (RDI) or alarm indication signal (AIS) alarms when the router is rebooted. [PR/257419: This issue has been resolved.]
- When both PIM and OSPF are configured on an IQ2 PIC, OSPF may lose adjacency if protocol PIM is removed. [PR/257848: This issue has been resolved.]
- After a graceful Routing Engine switchover on the M10i router, alarms might not resynchronize to the new primary Routing Engine. [PR/258034: This problem is resolved.] [PR/258034: This issue has been resolved.]
- On an MX-series router with a 4-port Gigabit Ethernet DPC, when you configure asynchronous notification, it does not function properly. [PR/259304: This issue has been resolved.]
- On the J2320 and J2350 series Services Routers, Gigabit Ethernet interfaces sometimes stop transmitting the packets if Transmit Descriptors are not updated with TransmitDone status properly. To work around this issue, restart pic 0. [PR/261010: This problem is resolved.]

### Services Applications

- When you configure twice NAT with static source and static destination translation, the destination port for ICMP flows might change (the ports are supposed to remain unchanged). [PR/96701: This issue has been resolved.]
- If Network Address Port Translation (NAPT) is configured and multiple short-lived flows are established, ports on AS PICs might not be assigned correctly. In some cases, this situation causes the AS PIC to stop functioning. [PR/229287: This issue has been resolved.]
- On the MultiServices 400 PIC, a memory warning flag might be set even with low traffic rates. [PR/251908: This issue has been resolved.]
- The `show services pgcp active-configuration` command does not display byte units for the "MG maximum PDU size" and "MGC maximum PDU size" output fields. [PR/256801: This issue has been resolved.]

### Routing Protocols

- Trace pointers for some BGP tasks were not updated on reconfiguration. [PR/69321: This issue has been resolved.]
- The `show pim source` does not display the correct information for both direct and non-direct sources. [PR/253629: This issue has been resolved.]
- PIM anycasts do not working when the source is connected to the Rendezvous Point router. [PR/256637: This issue has been resolved.]
- Route target filtering breaks when the last community received is withdrawn, because the route-filtering logic is being bypassed. Absence of any route target received from the peer is being treated as if 0/0 default was received from the peer. [PR/257011: This issue has been resolved.]

- The `show multicast snooping route bridge-domain name source-prefix prefix/length` command causes the multicast snooping process (`mcsnoopd`) to stop functioning. The workaround is not to use the `source-prefix` option. [PR/257788: This issue has been resolved.]
- The routing processes can generate nonfatal coredumps. [PR/258134: This issue has been resolved.]
- Deconfiguration of a routing instance on a router configured with uRPF may cause the routing process to restart. [PR/259727: This issue has been resolved.]

### MPLS Applications

- A router configured with a point-to-multipoint transmit-switch connection might stop functioning if the transmit label-switched path of the connection flaps. [PR/229175: This issue has been resolved.]
- You might encounter an interoperability issue when Cisco IOS-XR or IOS includes the `node-id` sub-object as part of the RRO in Reseveration messages. The JUNOS software is unable to find the next-next-hop router's interface address to signal a node-protecting bypass LSP. [PR/237491: This issue has been resolved.]
- MVPN P2MP deactivating of a vt interface in a particular VPN, say VPN-A on receiver PE, affects multicast traffic forwarding in other VPNs for a few seconds. This issue is also experienced during the activate sequence of the vt interface in VPN-A. [PR/252697: This issue has been resolved.]
- If an MPLS LSP configured with fast reroute is not advertised in to the IGP, that LSP might reuse an old unicast list and cause traffic drops. [PR/253352: This issue has been resolved.]
- If the authentication method for a LDP session is changed from using `authentication-key` to using `authentication-key-chain` or vice-versa, other unrelated LDP sessions may flap in addition to the affected LDP session flapping. [PR/258395: This issue has been resolved.]
- If there are many LSP flaps, used for the `p2mp-receive-switch`, and graceful restart is enabled, the interface might stop forwarding traffic. To recover, you need to deactivate and then activate the protocol connection `p2mp-receive-switch` configuration. As a workaround, disable graceful restart [PR/264930: This issue has been resolved.]

### VPNs

- When configuring multicast VPN for point to multipoint, the `tunnel-limit` statement for dynamic selective provider tunnels is not functioning. [PR/250701: This issue has been resolved.]
- A receiver directly attached over an Ethernet connection to a sender PE configured with multicast VPN over point to multipoint fails to receive multicast traffic. No such issue is observed when the receiver is connected over a Sonet or ATM link. [PR/252314: This issue has been resolved.]
- When you configure the `vlan-tags` statement under routing instances for a VPLS or a virtual switch instance, the JUNOS software might produce commit check

errors messages. The workaround is to use the `vlan-tags` statement in the configuration for interfaces only. [PR/256958: This issue has been resolved.]

### Class of Service

- On the M320 and T-series routing platforms, when you map multiple forwarding classes to the same queue and then include the multiple of those definitions in the scheduler map, the configuration might fail. [PR/103370: This issue has been resolved.]
- The behavior aggregate (BA) classifier is not applied to a logical interface configured with the encapsulation `ether-over-atm-llc` statement. [PR/255742: This issue has been resolved.]
- For IQ PICs on the M-series and T-series routers and Enhancing Queueing DPCs on the MX-series routers, when you configure a scheduler map with a queue configured with priority strict-high, in certain situations such as when a PIC is bounced, the incorrect queue buffer might be calculated. [PR/256263: This issue has been resolved.]
- When a PIC goes offline and then online, the following message might display on the Packet Forwarding Engine console: “cosman\_compute\_mad\_state: No ifd for ifd\_index < ifd index value > .” The message does not indicate any effect on the operation of the router unless a temporal or delay buffer configuration is present on the router. [PR/257814: This issue has been resolved.]
- If you configure a traffic shaper of over 75 MB on a Gigabit Ethernet interface, its overall throughput might decrease. For Gigabit Ethernet interfaces, configure shapers with a size of less than 75 MB. [PR/257951: This issue has been resolved.]

### Forwarding and Sampling

- If you configure a policer with a burst size limit larger than 67 MB, interfaces to which the policer is applied might not forward traffic. On some platforms the limit is higher, and the limit also depends on the available bandwidth. [PR/99758: This issue has been resolved.]
- Firewall filter source-address match sometimes does not function properly because of corner cases during firewall optimization. The workaround is to rearrange the term order of the filter. [PR/262491: This issue has been resolved.]

### Network Management

- When you configure an SNMP client list with a logical router, clients are not restricted. [PR/254574: This issue has been resolved.]

## 8.4R2

The following issues have been resolved since JUNOS Release 8.4R2. The identifier following the description is the tracking number in our bug database.

## Platform and Infrastructure

- Messages like the following might appear in the system log, indicating that the system clock went backward by 1 microsecond: "check\_kernel\_exec\_time: microuptime() went backwards (*seconds.microseconds* - > *seconds.microseconds*)."  
There is no operational effect. [PR/77411: This issue has been resolved.]
- On M20 routers, if you take an FPC offline, extract a PIC, reinsert the PIC, and attempt to bring the FPC online, the online operation might fail, the System and Switching Board (SSB) might dump core, and the router might reboot automatically. [PR/78988: This issue has been resolved.]
- Under some circumstances, the interface process (ifd) is interfering with the operation of an LSI interface. [PR/102431: This issue has been resolved.]
- If tricolor marking is configured on an interface and the interface is repeatedly disabled and enabled, memory that has been allocated might not be released properly. [PR/232472: This issue has been resolved.]
- If there are large numbers of routes and next hops, the MultiServices PIC might not be able to allocate enough memory to install all of them in to its internal database. [PR/235368: This issue has been resolved.]
- If a routing platform uses an aggregated Ethernet interface as the only internal (IGP) facing interface and has many external peering sessions on other interfaces not traversing the aggregated Ethernet interface, when the aggregated Ethernet interface goes down and up several times, the memory buffer of the routing platform might overflow. This can cause a switchover to the backup Routing Engine (on dual Routing Engine platforms) or cause the Routing Engine to become unreachable (on single Routing Engine platforms). [PR/236258: This issue has been resolved.]
- On MX-series Ethernet Services Routers with VPLS, when a core-facing interface that is configured for integrated routing and bridging (IRB) changes to be a regular routing interface, traffic might be discarded. As a workaround, restart the chassis process (chassisd). [PR/237212: This issue has been resolved.]
- If you configure a large number of MD5 authentication keys for BGP sessions, and then deactivate and reactivate the keys, the router might generate a commit error, and MD5 authentication might not be applied on some of the BGP sessions. [PR/237690: This issue has been resolved.]

## User Interface and Configuration

- If the configuration includes a commit script that uses the `jcs:invoke` routine, the router fails to boot successfully. [PR/95960: This issue has been resolved.]
- When you activate and deactivate VRF routing instances, the routing protocol process might generate a core file and stop operating. [PR/102088: This issue has been resolved.]
- In JUNOS Release 8.2 and later releases, user permissions are sometimes not calculated correctly. As a result, users might not be able to perform actions for which they have the required permission. [PR/229424: This issue has been resolved.]
- If the syslog host is configured and you commit a configuration change to the event process, an event policy execution causes the child event process to freeze.

A subsequent commit operation to the event process causes the main event process to fail. The child event process remains active. [PR/230064: This issue has been resolved.]

- If the event process (eventd) receives corrupted data from another process, it might generate a core file. [PR/236599: This issue has been resolved.]
- If a NETCONF client sends the <commit/> tag to a router on which **commit synchronize** is the default commit method, the <rpc-reply> tag element generated by the NETCONF server might not be well-formed XML. [PR/241659: This issue has been resolved.]

## Interfaces and Chassis

- On Channelized STM1 PICs, a tributary unit alarm indication signal (TU-AIS) alarm enabled for one channel might cause another channel to shut down. [PR/55357: This issue has been resolved.]
- On M160 and M40e routers, when you commit a configuration change, the router might generate a system log message that erroneously reports the master Packet Forwarding Engine Clock Generator (PCG) status as removed or offline. [PR/58716: This issue has been resolved.]
- When you manually disable the active Automatic Protection Switching (APS) interface to switch between the working and protect circuits, APS might not function properly. [PR/71083: This issue has been resolved.]
- With JUNOS Release 8.0 or later, some XENPAK transceivers might report spurious temperature, laser bias, laser output, and receive optical power alarms that are cleared again after about 10 seconds. There is no operational impact. [PR/98428: This issue has been resolved.]
- On J-series Services Routers, configuring more than one VRRP group on a port puts the port in to promiscuous mode. Forwarding performance can be affected, and duplicate ICMP messages might be sent in response to the **ping** command. This problem applies to 1-port Gigabit Ethernet ePIMs on all J-series platforms, and built-in Gigabit Ethernet interfaces on the J4350 and J6350 Services Routers. [PR/99796: This issue has been resolved.]
- When graceful Routing Engine switchover is configured and a switchover occurs, the kernel might reset and generate a core file. [PR/101359: This issue has been resolved.]
- On M320 routers, if the ingress and egress interfaces used for a TCC connection are installed in different types of FPCs (one in a standard FPC, another in an Enhanced III FPC), TCC encapsulation does not work. As a workaround, install the PICs in the same type of FPC. [PR/102997: This issue has been resolved.]
- When a switchover event occurs on a routing platform with a configuration that includes a large number of IPv6 routing instances, the Routing Engine that was previously the master might not be able to synchronize the kernel database. [PR/105268: This issue has been resolved.]
- If configuration for an interface is defined in both an internal system file and the JUNOS configuration, an internal database maintained by interface process (dcd) might include invalid pointers. When the process tries to free the pointers, it generates a core file. [PR/231145: This issue has been resolved.]

- When graceful restart is configured for routing protocols on a TX Matrix router, multicast routing might fail when a member link of an aggregated Ethernet interface is taken down. [PR/231772: This issue has been resolved.]
- On T640 and M320 routing platforms, if you configure link services IQ (LSQ) services and a link PIC to provide links to the LSQ bundles, when the FPC that houses the link PIC is taken offline, the `show interfaces` command might display invalid statistics for interfaces on the PIC. [PR/234521: This issue has been resolved.]
- The link fault management process might fail when you add an interface configuration to a deactivated hierarchy, and then roll back the configuration. [PR/234753: This issue has been resolved.]
- On a Channelized DS3 IQ PIC with loopback enabled for `ct1` or `t1` interfaces in a bundle, if you use the configuration mode `copy` command to create a new interface in the bundle based on one of the existing interfaces, the new interface does not come online until the PIC resets. [PR/235228: This issue has been resolved.]
- Adding or deleting a T1 member link causes one to two packets to drop on unrelated multilink PPP interfaces. [PR/236014: This issue has been resolved.]
- If you use an ATM2 Intelligent Queuing (IQ) interface as the protect circuit of an Automatic Protection Switching (APS) connection, and a graceful Routing Engine switchover (GRES) event occurs, the logical interface might be marked as hardware down rather than device down. If an APS circuit switchover happens after the GRES switchover, the ATM2 IQ interface does not clear the `aps_disable` flag, and traffic loss might occur on the connection. [PR/236610: This issue has been resolved.]
- When IPv6 route accounting is configured on a TX Matrix platform (the `route-accounting` statement is included at the `[edit forwarding options family inet6]` hierarchy level), the configuration is not copied to the T640 routing nodes in the matrix. As a result, the fields in the `IPv6 transit statistics` section of the output from the `show interfaces` command do not increment. [PR/237054: This issue has been resolved.]
- When a graceful Routing Engine switchover occurs and a member link of an aggregated Ethernet interface is taken offline, multicast traffic might not be rerouted through other member links as expected. [PR/237098: This issue has been resolved.]
- If you apply more than one class-of-service rewrite rule to interfaces on an IQ2 PIC, only the first rule works correctly on all interfaces. [PR/238250: This issue has been resolved.]

### Services Applications

- If you perform numerous commit operations, the key management process (kmd) might generate a core file and stop operating because of an incorrect memory allocation procedure. [PR/232085: This issue has been resolved.]
- When multiple Layer 2 Tunneling Protocol (L2TP) tunnels from different tunnel groups are being set up simultaneously, the same tunnel identifier might be assigned to more than one tunnel, and some tunnel attributes might be assigned

to the wrong L2TP session. As a workaround, issue the `clear services l2tp tunnel` command to clear the L2TP tunnels. [PR/233184: This issue has been resolved.]

- If the configuration includes certain features (for example, firewalls), service PICs do not come online if it takes the Packet Forwarding Engine more than 2 minutes to establish communication between PICs and the Routing Engine. [PR/236926: This issue has been resolved.]
- When two or more IQ2 PICs are installed on an M7i, M10i, or M120 router, Layer 2 Tunneling Protocol (L2TP) policer sessions are not established correctly. As a workaround, issue the `request chassis pic fpc-slot slot-number pic-slot slot-number` command to take all but one of the PICs offline. [PR/237356: This issue has been resolved.]

## Routing Protocols

- Label aggregation is not working properly when route reflection is also configured in Layer 3 VPN networks. [PR/228039: This issue has been resolved.]
- When you use the `clear bgp damping address` command to remove a damped route from a routing table and specify a prefix on the address (for example, `clear bgp damping 10.0.0.8/30`), the route might not be removed. (Damped routes are listed as `hidden` in the output of the `show route damping suppressed terse` command.) [PR/231502: This issue has been resolved.]
- Configuring a PIM bootstrap export policy (including the `export` statement at the `[edit protocols pim rp bootstrap family family-name]` hierarchy level) has no effect. PIM bootstrap import policies work as expected. [PR/232200: This issue has been resolved.]
- If a local route is imported from another routing instance, it might not be installed in to the forwarding table. A message like the following might appear in the system log: `"timestamp router rpd[PID]: KRT ADD for 10.10.10.10/32 = > { ifl interface-index addr 10.10.10.10 } failed, error "ENOENT -- Item not found"."` [PR/234918: This issue has been resolved.]
- The Route Distinguisher field in the output of the `show route receive-protocol` command reports the value configured on the local router (by the `route-distinguisher` statement at the `[edit routing-instances routing-instance-name]` hierarchy level), instead of the route-distinguisher value configured on the remote PE router. [PR/239698: This issue has been resolved.]

## MPLS Applications

- If the configuration for an LSP includes more than 16 equal-cost multipaths (ECMPs), RSVP might not function correctly. [PR/228338: This issue has been resolved.]
- RSVP PATH messages are not bundled when RSVP aggregation is negotiated on a point-to-point link. [PR/234545: This issue has been resolved.]
- The `mplsLspPathUp` and `mplsLspPathDown` SNMP traps are enabled when either or both of the `(syslog | no-syslog)` and `(trap | no-trap)` statements are included at the `[edit protocols mpls log-updown]` hierarchy level. The correct behavior is for the traps to be enabled only when the `trap-path-down` and `trap-path-up` statements

are included at the `[edit protocols mpls log-updown]` hierarchy level. [PR/237464: This issue has been resolved.]

## VPNs

- If you configure a Layer 2 circuit across a logical tunnel interface that uses `ethernet-vpls` encapsulation, the Layer 2 circuit connection might not come up. As a workaround, use `ethernet` encapsulation on the logical tunnel interface. [PR/100161: This issue has been resolved.]
- On a provider edge (PE) router configured with multiprotocol BGP-based multicast VPNs and connected directly to a receiver, if you modify the multicast VPN import target with the `import-target` statement at the `[edit routing-instances routing-instance-name mvpn route-target]` hierarchy level, the BGP route reflector might fail to readvertise the multicast VPN routes. As a workaround, issue the `clear bgp neighbor soft` command on the route reflector to force it to readvertise all the multicast VPN routes without resetting the BGP sessions. [PR/104192: This issue has been resolved.]
- On MX-series Ethernet Services Routers, if the VPLS configuration includes the `no-tunnel-services` statement (meaning that an LSI is used instead of a virtual tunnel), customer edge (CE)-facing interfaces sometimes discard traffic. In the output from the `show interfaces` command, the list in the **Flags** field for the logical interface might include the CCC-Down or Hardware-Down flag. As a workaround, deactivate and reactivate the CE interface. [PR/234756: This issue has been resolved.]
- For MX-series routers, if you change the route distinguisher used for a VPLS routing instance, VPLS connections might not come up. As a workaround, deactivate and then reactivate the VPLS routing instance. [PR/235048: This issue has been resolved.]

## Class of Service

- When the combination of interface types configured on a channelized IQ PIC requires an amount of memory that exceeds the maximum available, some of the interfaces might not handle traffic correctly. [PR/102932: This issue has been resolved.]
- Some JUNOS processes, including the class-of-service process (`cosd`), might not correctly release memory that they have allocated. [PR/230771: This issue has been resolved.]
- A class-of-service EXP classifier applied to a VRF routing instance does not work correctly for traffic received by a PIC housed in an Enhanced Scaling FPC (FPC4). [PR/233694: This issue has been resolved.]
- If you configure 600 T1 links and 255 Multilink Frame Relay (FRF.16) bundles on link services queuing (LSQ) interfaces in an AS PIC or MultiServices PIC, then downgrade the routing platform to JUNOS Release 8.4R1 from a later release, the downgrade might fail the validation process and generate the following error message: “Current configuration not compatible with `/var/tmp/jinstall-8.4R1.13-domestic-signed.tgz`”. As a workaround, deactivate the class-of-service configuration, downgrade to JUNOS Release 8.4R1, and then

reactivate the class-of-service configuration. [PR/234031: This issue has been resolved.]

- Over T3 or E3 links on J-series Services Routers, traffic that is assigned strict-high priority might be discarded when large-packet high-priority traffic is sent together with small-packet low-priority traffic, and the high-priority queue's buffer size is much lower than that of the low-priority queue. As a workaround, configure a shaper with its rate set to match the line rate. [PR/234626: This issue has been resolved.]
- Deactivating and then activating an interface causes packet loss on other ports of the same PIC. [PR/238314: This issue has been resolved.]

### Forwarding and Sampling

- On an M120 router, if the `then` statement in a firewall filter (at the `[edit firewall filter filter-name term term-name]` hierarchy level) includes both a `policer` statement and a `count` statement, the filter might not handle packets as specified by the configuration. [PR/105465: This issue has been resolved.]
- On MX-series or M120 routers, if you use a tunnel services-based PIC to configure multiple VPLS instances over virtual loopback tunnel (VT) interfaces, and then include the `no-tunnel-services` statement to direct VPLS traffic over label-switched interface (LSI) logical interfaces instead of the VT interfaces, in some cases, the VPLS traffic might stop even though the VPLS connections remain up. As a workaround, deactivate and reactivate interfaces in the affected VPLS instances. [PR/228411: This issue has been resolved.]
- If you apply the same firewall to both IPv4 and IPv6 traffic on an interface by including the `input-list` statement at the `[edit interfaces interface-name unit logical-unit-number family inet filter]` and `[edit interfaces interface-name unit logical-unit-number family inet6 filter]` hierarchy levels, the `clear firewall filter interface-name` command clears the counters for IPv4 bytes and packets but not for IPv6 bytes and packets (in the output from the `show firewall filter interface-name` command). [PR/229016: This issue has been resolved.]

### Network Management

- Bridge domain instances are not listed in the `VacmContextTable` table after a Routing Engine switchover or restart of the SNMP process. [PR/233765: This issue has been resolved.]

## Errata

---

This section lists outstanding issues with the documentation.

### Platform and Infrastructure

- The following note has been added to the section titled “Requirements for Routers with a Backup Router Configuration” in Chapter 6 of the *JUNOS 8.5 High Availability Configuration Guide*:



**NOTE:** If you have a backup router configuration in which multiple static routes point to a gateway from `fxp0`, you must configure prefixes that are more specific than the static routes or include the `retain` flag at the `[edit routing-options static route]` hierarchy level.

For example, if you configure the static route `172.16.0.0/12` from `fxp0` for management purposes, you must specify the backup router configuration as follows:

```
backup-router 172.29.201.62 destination [172.16.0.0/13 172.16.128.0/13]
```

---

*[High Availability]*

- The following errata replaces the entire section titled “Graceful Routing Engine Switchover System Requirements” in Chapter 5 of the *JUNOS 8.5 High Availability Configuration Guide*.

Graceful Routing Engine switchover supports most Physical Interface Cards (PICs) on the M10i, M20, M40e, M120, M320, T320, T640TX Matrix, and T1600 routers with the appropriate version of JUNOS software. Use Table 1 on page 73 to determine which JUNOS software release provides initial graceful Routing Engine switchover support for PICs installed in your router. A dash (–) indicates that support is not available. Except where noted, the JUNOS software release is R1.

---



**NOTE:** When an unsupported PIC is online, you cannot enable graceful Routing Engine switchover. If graceful Routing Engine switchover is already enabled, an unsupported PIC cannot come online.

---

The following constraints apply to graceful Routing Engine switchover PIC support:

- You can include the `graceful-switchover` statement at the `[edit chassis redundancy]` hierarchy level on a router with Adaptive Services and MultiServices PICs configured on it and issue the `commit` command. The `commit` succeeds. However, all services on these PICs are reset during a switchover.
- Graceful Routing Engine switchover is not supported on any Monitoring Services PICs or Multilink Services PICs. If you include the `graceful-switchover` statement at the `[edit chassis redundancy]` hierarchy level on a router with either of these PIC types configured on it and issue the `commit` command, the `commit` fails.
- Graceful Routing Engine switchover is not supported on MultiServices 400 PICs configured for monitoring services applications. If you include the `graceful-switchover` statement, the `commit` fails.

**Table 1: Graceful Routing Engine Switchover PIC Support**

PIC Type	M10i	M20	M40e	M120	M320	T320	T640	TX Matrix	T1600
Adaptive Services PICs	7.3	7.3	7.3	–	–	–	–	–	–
Adaptive Services II PICs	7.2	7.2	7.2	8.2	7.2	7.2	7.2	7.3	8.5
Adaptive Services II PICs with Link Services IQ (LSQ) interfaces	7.6	7.6	7.6	8.2	7.6	7.6	7.6	–	8.5
Adaptive Services II FIPS PICs	7.3	–	7.3	8.2	7.3	–	–	–	–
ATM2 DS3 IQ PICs	6.1	6.1	6.1	8.2	6.4	–	8.0	8.0	8.5
ATM2 E3 IQ PICs	6.1	6.1	6.1	8.2	6.3	7.4	7.4	7.4	8.5
ATM2 OC3/STM1 IQ PICs	6.1	6.0	6.0	8.2	6.2	6.0	7.6	7.6	8.5
ATM2 OC12/STM4 IQ PICs, 1-port	6.1	6.0	6.0	8.2	6.2	6.0	8.0	8.0	8.5
ATM2 OC12/STM4 IQ PICs, 2-port	–	–	6.0	8.2	6.2	6.0	6.0	7.0	8.5
ATM2 OC48/STM16 IQ PICs with SFP	–	–	7.3	8.2	7.3	7.3	7.3	7.3	8.5
Channelized DS3 PICs	6.1	6.0	6.0	–	–	–	–	–	–
Channelized OC12 PICs	6.1	6.0	6.0	–	–	–	–	–	–
Channelized DS3 IQ PICs	6.1	6.1	6.1	8.2	6.2	6.3	8.0	8.0	8.5
Channelized E1 IQ PICs	6.1	6.1	6.1	8.2	6.2	–	–	–	–
Channelized OC3 IQ PICs	7.1	7.1	7.1	8.2	7.1	7.1	7.6	7.6	8.5

**Table 1: Graceful Routing Engine Switchover PIC Support** (continued)

PIC Type	M10i	M20	M40e	M120	M320	T320	T640	TX Matrix	T1600
Channelized OC12 IQ PICs	6.1	6.1	6.1	8.2	6.2	6.1	6.3	7.0	8.5
Channelized STM1 IQ PICs	6.1	6.1	6.1	8.2	6.2	6.1	7.5	7.5	8.5
Channelized T1 IQ PICs	7.4	7.4	7.4	8.2	7.4	–	–	–	–
DS3 PICs	6.1	6.0	6.0	8.2	6.2	6.0	6.3	7.0	8.5
E1 PICs	6.1	6.0	6.0	8.2	6.4	–	–	–	–
E3 PICs	6.1	–	–	–	6.3	–	–	–	–
E3 IQ PICs	6.1	6.1	6.1	8.2	6.2	6.2	6.3	7.3	8.5
EIA-530 PICs	6.1	6.0	6.0	–	–	–	–	–	–
ES PICs	6.1	6.1	6.1	–	6.2	6.1	–	–	–
Fast Ethernet PICs, 4-port	6.1	6.0	6.0	8.2	6.2	6.0	6.3	7.0	8.5
Fast Ethernet PICs, 8-port	6.1	6.0	6.0	8.2	6.3	–	–	–	–
Fast Ethernet PICs, 12-port	6.1	6.0	6.0	8.2	6.2	6.1	–	–	–
Fast Ethernet PICs, 48-port	–	–	6.0	8.2	6.4	–	–	–	–
Fast Ethernet PICs (aggregated)	6.2	6.2	6.2	8.2	6.2	6.2	6.2	7.0	8.5
Gigabit Ethernet PICs with SFP, 1-port	6.4	6.4	6.4	8.2	6.4	6.4	8.0	8.0	8.5
Gigabit Ethernet PICs with SFP, 1-port	6.4	6.4	6.4	8.2	6.4	6.4	8.0	8.0	8.5
Gigabit Ethernet PICs with SFP, 2-port	–	–	6.4	8.2	6.4	6.4	6.4	7.0	8.5
Gigabit Ethernet PICs with SFP, 4-port	–	–	7.0	8.2	7.0	7.0	7.0	7.0	8.5

**Table 1: Graceful Routing Engine Switchover PIC Support** (continued)

PIC Type	M10i	M20	M40e	M120	M320	T320	T640	TX Matrix	T1600
Gigabit Ethernet PICs with SFP, 10-port	–	–	–	8.2	6.2	6.0	6.0	7.0	8.5
Gigabit Ethernet IQ PICs, 1-port	7.0	7.0	7.0	8.2	7.0	7.0	8.0	8.0	8.5
Gigabit Ethernet IQ PICs, 2-port	–	–	7.0	8.2	7.0	7.0	7.0	7.3	8.5
Gigabit Ethernet IQ2 PICs, 4-port	7.6	–	7.6	8.2	7.6	7.6	7.6	7.6	8.5
Gigabit Ethernet IQ2 PICs, 8-port	–	–	7.6	8.2	7.6	7.6	7.6	7.6	8.5
10-Gigabit Ethernet DWDM PICs	–	–	–	8.2	7.5	7.5	7.5	7.5	8.5
10-Gigabit Ethernet XENPACK PICs	–	–	–	8.2	6.2	6.2	6.2	7.0	8.5
Link Services PICs	7.0	7.0	7.0	–	7.0	–	–	–	–
MultiServices 100	8.1R2	–	8.1R2	8.2	8.1R2	8.1R2	8.1R2	8.1R2	8.5
MultiServices 400	–	–	8.1R2	8.2	8.1R2	8.1R2	8.1R2	8.1R2	8.5
MultiServices 500	–	–	–	8.3	8.3	8.3	8.3	8.3	8.5
SONET/SDH OC3c/STM1 PICs, 2-port Type 1	6.1	–	–	–	–	–	–	–	–
SONET/SDH OC3c/STM1 PICs, 4-port Type 1	6.1	6.0	6.0	8.2	6.2	7.1	7.1	7.1	8.5
SONET/SDH OC3c/STM1 PICs, 4-port Type 2	–	–	–	–	–	6.0	6.0	7.0	8.5

**Table 1: Graceful Routing Engine Switchover PIC Support** (continued)

PIC Type	M10i	M20	M40e	M120	M320	T320	T640	TX Matrix	T1600
SONET/SDH OC3/STM1 PICs with SFP, 2-port Type 1	8.4	8.4	-	-	-	-	-	-	-
SONET/SDH OC3/STM1 (Multi-Rate) with SFP, 4-port Type 1	8.4	8.4	8.4	8.4	8.4	8.4	8.4	8.4	8.5
SONET/SDH OC3/STM1 (Multi-Rate) with SFP, 4-port Type 2	-	-	8.3	8.3	8.3	8.3	8.3	8.3	8.5
SONET/SDH OC12c/STM4 SMIR PICs, 1-port Type 1	6.1	6.0	6.0	8.2	6.2	7.5	6.0	7.0	8.5
SONET/SDH OC12c/STM4 MM PICs, 1-port Type 1	-	-	-	-	-	-	7.1	7.1	8.5
SONET/SDH OC12c/STM4 PICs, 4-port Type 2	-	-	6.0	8.2	6.2	6.0	6.0	7.0	8.5
SONET/SDH OC12/STM4 (Multi-Rate) with SFP, 1-port Type 1	8.4	8.4	8.4	8.4	8.4	8.4	8.4	8.4	8.5
SONET/SDH OC12/STM4 (Multi-Rate) with SFP, 4-port Type 2	-	-	8.3	8.3	8.3	8.3	8.3	8.3	8.5
SONET/SDH OC48c/STM16 PICs with SFP, 1-port quad-wide Type 1	6.4	6.1	-	-	-	-	-	-	-
SONET/SDH OC48c/STM16 PICs with SFP, 1-port Type 2	-	-	6.1	8.2	6.2	6.1	6.1	7.0	8.5

**Table 1: Graceful Routing Engine Switchover PIC Support** (continued)

PIC Type	M10i	M20	M40e	M120	M320	T320	T640	TX Matrix	T1600
SONET/SDH OC48c/STM16 PICs with SFP, 4-port Type 3	–	–	–	8.2	6.2	6.2	6.2	7.0	8.5
SONET/SDH OC48/STM16 (Multi-Rate) with SFP, 1-port, Type 2	–	–	8.3	8.3	8.3	8.3	8.3	8.3	8.5
SONET/SDH OC192c/STM64 PICs, 1-port Type 3	–	–	–	8.2	6.2	6.0	6.0	7.0	8.5
SONET/SDH OC768c/STM256 PICs, 1-port Type 4	–	–	–	–	–	–	7.5	7.5	8.5
SONET/SDH PICs (aggregated)	6.2	6.2	6.2	8.2	6.2	6.2	6.2	7.0	8.5
T1 PICs	6.1	6.0	6.0	8.2	6.4	–	–	–	–
Tunnel Services PICs, Type 1	6.1	6.0	7.0	8.2	6.2	6.1	8.0	8.0	8.5
Tunnel Services PICs, Type 2	–	–	7.0	8.2	6.2	6.1	6.1	7.0	8.5
Tunnel Services PICs, Type 3	–	–	–	8.2	6.3	6.0	6.0	7.0	8.5
Tunnel Services PICs, 40-gigabit	–	–	–	–	–	–	8.0	8.0	8.5

[High Availability]

### User Interface and Configuration

- When you configure a logical-router system administrator, you cannot also configure graceful Routing Engine switchover (GRES) on the router. [PR/237070]
- J-Web Quick Configuration pages do not support IPv6 addressing and routing. [J-series Basic Configuration]

- The new `bridge` option to the `show system statistics` command displays system statistics on MX-series routers. The option is not documented in the *JUNOS System Basics and Services Command Reference*.
- The new `static-mac` statement at the `[edit routing-instances instance-name protocols l2vpn site site-name interface interface-name]` and `[edit routing-instances instance-name protocols vpls site site-name interface interface-name]` is not documented in the *JUNOS VPNs Configuration Guide*.

### Interfaces and Chassis

- To display the FRU model number, part number, and serial number, issue the `show chassis hardware models` command. To display the FRU model number, part number, and CLEI code, issue the `show chassis hardware clei-models` command. [*System Basics and Services Command Reference*]
- FRF.12 is supported on link services (ls-) interfaces on the J-series routing platform. [*Services Interfaces*]
- The drop-and-insert multiplexer is now integrated in to channelized T1/E1 PIMs on J-series Services Routers. The `data-input (system | interface interface-name)` statement at the `[edit interfaces ds-pim/O/port:channel]` hierarchy level is not documented in the *JUNOS Network Interfaces Configuration Guide*.
- To configure the Link Aggregation Control Protocol (LACP), include the `lACP` statement at the `[edit protocols]` hierarchy level. [*Network Interfaces Configuration Guide*]
- **IP address for the master Routing Engine**—Enables you to configure a management IP address that is always used by the master Routing Engine, even when a failover occurs. To configure, include the `master-only` statement at the `[edit interfaces fxp0 unit logical-unit-number inet address ip-address]` hierarchy level on the master Routing Engine. [*Network Interfaces Configuration Guide*]

### Services Applications

- The documentation does not include information about the `clear-ike-sas-on-pic-restart` statement, which you can include at the `[edit services ipsec-vpn]` hierarchy level to have the IKE SAs associated with tunnels cleared automatically when an AS PIC goes offline or is restarted. [*Services Interfaces*]

### General Routing

- The manuals currently state that only the following routing platforms support GRES for VPLS: M10i, M20, M40e, M320, T320, and T640. The TX Matrix routing platform also supports GRES for VPLS. [*System Basics, Feature Guide, VPNs*]

### Routing Protocols

- For the Spanning Tree Protocol (STP) bridge domain configuration and `show spanning-tree bridge` command shown in the *Routing Protocols Configuration Guide*, the words `bridge domain` should be `bridge`. In addition, the extended system ID is now configurable. [*Routing Protocols Configuration Guide*]

- To configure Layer 2 control protocol features, include the `layer2-control` statement at the `[edit protocols]` hierarchy level. *Routing Protocols Configuration Guide*

### MPLS Applications

- Inter-AS traffic engineering (Phase 2) enables traffic-engineered MPLS LSPs to dynamically discover OSPF autonomous system boundary routers (ASBRs) and enables routers to establish a traffic-engineered LSP across multiple autonomous systems (ASs). Each AS is assumed to be under the control of a single service provider and to use OSPF. To configure traffic engineering across multiple ASs using OSPF, include the `traffic-engineering` statement at the `[edit protocols (ospf | ospf3) area area-id interface interface-name passive]` hierarchy level. *[MPLS Applications]*
- The recommended range for static MPLS label-switched path (LSP) labels is now 1,000,000 through 1,048,575. The previous recommended range for static labels was 10,000 through 99,999. Configurations that have static labels outside the new range cannot be committed. To configure static MPLS labels, include the `label-map` statement at the `[edit protocols mpls interface interface-name]` hierarchy level. The JUNOS software now uses the range of 10,000 through 99,999 for processing label-switched interface (LSI) labels. *[MPLS Applications]*

### VPNs

- The statement hierarchy for the `vpls` statement shown in the “Configuring the VPLS Routing Instance” section includes an extraneous bracket character (`}`). The corrected statement hierarchy is as follows:

```
vpls {
  active-interface {
    any;
    primary interface-name;
  }
  interface-mac-limit limit;
  mac-table-size size;
  neighbor neighbor-id;
  no-tunnel-services;
  site site-name {
    active-interface {
      any;
      primary interface-name;
    }
    interface interface-name {
      interface-mac-limit limit;
    }
    multi-homing;
    site-identifier identifier;
    site-preference preference-value;
  }
  site-range number;
  traceoptions {
    file filename <replace> <size size> <files number> <no-stamp>;
    flag flag <flag-modifier> <disable>;
  }
}
```

```

tunnel-services {
    devices device-names;
    primary primary-device-name;
}
vpls-id vpls-id;
}

```

[VPNs]

- You cannot configure the `mac-table-aging-time` statement on MX-series routers. [VPNs]

### Class of Service

- As stated in the documentation, you can selectively set the DSCP field of IPv4 and IPv6 packets to 0 without affecting output queue assignment, and continue to set the MPLS EXP field according to the configured rewrite table, based on forwarding classes. This feature is not supported with GRE and IP-IP tunnels. The documentation incorrectly implies that you can use the `dscp 0` action modifier to set the DSCP field of IPv4 and IPv6 packets to 0. For IPv4 traffic, the `dscp 0Service` action modifier at the `[edit firewall family inet filter filter-name term term-name then]` hierarchy level is valid. However, for IPv6 traffic, you configure this feature by including the `traffic-class 0` action modifier at the `[edit firewall family inet6 filter filter-name term term-name then]` hierarchy level. [CoS]
- On J-series Services Routers, if you create a policy to match IPv4 traffic with a route filter, assign the traffic to a forwarding class, and then apply the policy at the `[edit class-of-service forwarding-policy class policy-name]` hierarchy level, use of the `classification-override` statement at the `[edit class-of-service forwarding-policy class policy-name]` hierarchy level is not supported. [*J-series Advanced Configuration, CoS*]
- For 4-port Fast Ethernet ePIMs on J-series Services Routers, if you apply a CoS scheduler map on outgoing (egress) traffic, the router does not divide the bandwidth appropriately among the CoS queues. As a workaround, configure enforced CoS shaping on the ports. [*J-series Getting Started*]

### Routing Policy and Firewall Filters

- Floating point metric multiplier values in routing policies are now limited to eight significant digits. [*Policy*]
- In the documentation for the extended DHCP relay agent feature in the *JUNOS 8.3 Policy Framework Configuration Guide*, the syntax for the `server-group` statement is incorrectly documented as `server-groups`. You can include the `server-group` statement at the `[edit forwarding-options dhcp-relay]` hierarchy level to configure a named group of DHCP server addresses for use by the extended DHCP relay agent on the router. The correct syntax for the `server-group` statement is as follows:

```

[edit forwarding-options dhcp-relay]
server-group {
    server-group-name {
        server-ip-address;
    }
}

```

}

**JUNOS XML API and Scripting**

- In the *JUNOS 8.4 Configuration and Diagnostic Automation Guide*, which also applies to JUNOS Release 8.5, the section titled “Using Regular Expressions to Refine the Set of Events That Cause a Policy to Be Executed” includes a cross reference to a configuration example, but the target link is missing. The correct target is the section titled “Controlling Event Policy Using a Regular Expression” (<http://www.juniper.net/techpubs/software/junos/junos84/swconfig84-automation/controlling-event-policy-using-a-regular-expression.html>). [*Automation*]

**M-series, MX-series, and T-series Upgrade and Downgrade Instructions**

---

In JUNOS Release 8.5, the JUNOS software was extended to use FreeBSD version 6.1. As a result, the following requirements apply when you upgrade your routing platform to JUNOS Release 8.5 and later:

- For J-series, M-series, MX-series, and T-series routing platforms, the minimum requirement for installation media (such as a compact flash disk, internal flash disk, or PC Card) is 256 MB.
- For M-series, MX-series, and T-series routing platforms, you must perform the upgrade using the `jinstall` package.
- For all routing platforms, when upgrading from JUNOS Release 8.2 or below, you must use the `no-validate` option when you issue the `request system software add` command to perform the upgrade.
- For J-series Services Routers, you must perform the upgrade using the CLI. Do not use the Quick Configuration upgrade option in the J-Web user interface.

This section discusses the following topics:

- Upgrade to Release 8.5 on page 81
- Downgrade from Release 8.5 on page 84

**Upgrade to Release 8.5**

When upgrading or downgrading the JUNOS software, always use the `jinstall` package. Use other packages (such as the `jbundle` package) only when so instructed by a Juniper Networks support representative. For information about the contents of the `jinstall` package and details of the installation process, see the *JUNOS System Basics Configuration Guide*.



**NOTE:** Before upgrading, back up the file system and the currently active JUNOS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:

```
user@host> request system snapshot
```

The installation process rebuilds the file system and completely reinstalls the JUNOS software. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files) may be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the *JUNOS System Basics Configuration Guide*.

---

## **Upgrade Policy for JUNOS Software Extended End Of Life Releases**

---

A direct upgrade and downgrade path is now available for JUNOS Software Extended End of Life (EEOL) releases. You can upgrade directly from one EEOL release to the next release even though EEOL releases frequently occur in increments beyond three releases. The current upgrade and downgrade policy for a non-EEOL release is that you can only upgrade and downgrade by up to three releases at a time. The +3 policy remains unchanged for non-EEOL releases but includes a direct upgrade and downgrade path for EEOL to next EEOL releases

It is important to note that you can only upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release. For example, JUNOS Software Releases 8.5, 9.3, and 10.0 are EEOL releases. You can only upgrade from JUNOS Software release 8.5 to JUNOS Software Release 10.0 by first upgrading to JUNOS Software Release 9.3. This policy also applies to downgrades where you cannot skip an EEOL release but must target the EEOL release occurring directly before the currently installed EEOL release.

For more information on EEOL releases and to review a list of EEOL releases, see <http://www.juniper.net/support/eol/junos.html>

---

The download and installation process for JUNOS Release 8.5R4 is the same as for previous JUNOS releases.

If you are not familiar with the download and installation process, follow these steps:

1. Using a Web browser, follow the links to the download URL on the Juniper Networks Web page. Choose either **Canada and U.S. Version** or **Worldwide Version**:
  - <https://www.juniper.net/support/csc/swdist-domestic/> (customers in the United States and Canada)
  - <https://www.juniper.net/support/csc/swdist-ww/> (all other customers)
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Download the software to a local host.
4. Copy the software to the routing platform or to your internal software distribution site.
5. Install the new `jinstall` package on the routing platform.



**NOTE:** We recommend that you upgrade all software packages out-of-band using the console because in-band connections are lost during the upgrade process.

---

Customers in the United States and Canada use the following command:

```
user@host> request system software add validate reboot
source/jinstall-8.5R4.3-domestic-signed.tgz
```

All other customers use the following command:

```
user@host> request system software add validate reboot
source/jinstall-8.5R4.3-export-signed.tgz
```

Replace `source` with one of the following values:

- `/pathname`—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - `ftp://hostname/pathname`
  - `http://hostname/pathname`
  - `scp://hostname/pathname` (available only for Canada and U.S. version)

The `validate` option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Adding the `reboot` command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.



**NOTE:** After you install a JUNOS 8.5 `jinstall` package, you cannot issue the `request system software rollback` command to return to the previously installed software. Instead you must issue the `request system software add validate` command and specify the `jinstall` package that corresponds to the previously installed software.

---

### Downgrade from Release 8.5

To downgrade from Release 8.5 to another supported release, follow the procedure for upgrading, but replace the 8.5 `jinstall` package with one that corresponds to the appropriate release.



**NOTE:** You cannot downgrade more than three releases. For example, if your routing platform is running JUNOS Release 7.5, you can downgrade the software to Release 7.2 directly, but not to Release 7.1; as a workaround, you can first downgrade to Release 7.2 and then downgrade to Release 7.1.

---

For more information, see the *JUNOS System Basics Configuration Guide*.

### J-series Upgrade and Downgrade Instructions

---

In JUNOS Release 8.5, the JUNOS software was extended to use FreeBSD version 6.1. As a result, the following requirements apply when you upgrade your router to JUNOS Release 8.5 and later:

- To upgrade with the JUNOS CLI, the minimum requirement for installation media (such as a compact flash disk, internal flash disk, or PC card) is 256 MB. To use the J-Web interface for an upgrade, you must have 512 MB or more.
- For J-series Services Routers with a 256-MB compact flash:
  - You must perform the upgrade with the CLI. Do not use the J-Web interface for the upgrade.
  - Before upgrading to this release, see the important information in “Special Instructions for J-series Routers with a 256-MB Compact Flash” on page 92.
- When upgrading from JUNOS Release 8.2 or earlier, upgrade to an interim JUNOS Release 8.3 or later first. (Alternatively, you can use the `no-validate` option with the `request system software add` command, but we do not recommend this upgrade method.)

If the router is running a software version earlier than JUNOS Release 7.2R3 or 7.3R2, you might need to upgrade to one of these interim software releases before you can upgrade to JUNOS Release 8.3 or later.

This section contains the following topics:

- Upgrade and Downgrade Overview on page 85
- Before You Begin on page 86
- Downloading Software Upgrades from Juniper Networks on page 87
- Installing Software Upgrades with the J-Web Interface on page 87
- Installing Software Upgrades with the CLI on page 88
- Downgrade Instructions on page 90
- Special Instructions for J-series Routers with a 256-MB Compact Flash on page 92
- Cleaning Up Files on page 92
- Verifying Available Compact Flash Space on page 93
- Increasing the Compact Flash Space on page 93

## **Upgrade and Downgrade Overview**

Typically, you upgrade the JUNOS software on a Services Router by downloading a set of images onto your router or onto another system on your local network, such as a PC. You then uncompress the package and install the uncompressed software using the CLI. Finally, you boot your system with this upgraded device.

A JUNOS software package is a collection of files that make up a software component. You can download software packages either for upgrading JUNOS software or for recovering a primary compact flash.

All JUNOS software is delivered in signed packages that contain digital signatures, Secure Hash Algorithm (SHA-1) checksums, and Message Digest 5 (MD5) checksums. For more information about JUNOS software packages, see the *JUNOS Software Installation and Upgrade Guide*.

### **Upgrade Software Packages**

Download an upgrade software package, also known as an install package, to install new features and software fixes as they become available.

An upgrade software package name is in the following format:

*package-name-m.nZx-distribution.tgz*.

- *package-name* is the name of the package—for example, *junos-jseries*.
- *m.n* is the software release, with *m* representing the major release number—for example, 8.0.
- *Z* indicates the type of software release. For example, *R* indicates released software, and *B* indicates beta-level software.

- *x* represents the version of the major software release—for example, 2.
- *distribution* indicates the area for which the software package is provided—**domestic** for the United States and Canada and **export** for worldwide distribution.

A sample J-series upgrade software package name is `junos-jseries-8.0R2-domestic.tgz`.

## Recovery Software Packages

Download a recovery software package, also known as an install media package, to recover a primary compact flash device.

A recovery software package name is in the following format:

*package-name-m.nZx-export-cfnnn.gz*.

- *package-name* is the name of the package—for example, `junos-jseries`.
- *m.n* is the software release, with *m* representing the major release number—for example, 8.0.
- *Z* indicates the type of software release. For example, **R** indicates released software, and **B** indicates beta-level software.
- *x* represents the version of the major software release—for example, 2.
- **export** indicates that the recovery software package is the exported worldwide software package version.
- *cfnnn* indicates the size of the target compact flash device in megabytes—for example, `cf256`.

A sample J-series recovery software package name is `junos-jseries-8.0R2-export-cf256.gz`.

## Before You Begin

Before upgrading, be sure to back up the currently running and active file system and configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. To back up the file system, you must have a removable compact flash disk installed on a J4300 or J6300 Services Router, or a USB drive installed on any J-series Services Router. The backup device must have a storage capacity of at least 256 MB.

To back up the file system to the removable compact flash disk, issue the following command:

```
user@host> request system snapshot media removable-compact-flash
```

To back up the file system to the removable USB drive, issue the following command:

```
user@host> request system snapshot media usb
```

## **Upgrade Policy for JUNOS Software Extended End Of Life Releases**

---

A direct upgrade and downgrade path is now available for JUNOS Software Extended End of Life (EEOL) releases. You can upgrade directly from one EEOL release to the next release even though EEOL releases frequently occur in increments beyond three releases. The current upgrade and downgrade policy for a non-EEOL release is that you can only upgrade and downgrade by up to three releases at a time. The +3 policy remains unchanged for non-EEOL releases but includes a direct upgrade and downgrade path for EEOL to next EEOL releases

It is important to note that you can only upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release. For example, JUNOS Software Releases 8.5, 9.3, and 10.0 are EEOL releases. You can only upgrade from JUNOS Software release 8.5 to JUNOS Software Release 10.0 by first upgrading to JUNOS Software Release 9.3. This policy also applies to downgrades where you cannot skip an EEOL release but must target the EEOL release occurring directly before the currently installed EEOL release.

For more information on EEOL releases and to review a list of EEOL releases, see <http://www.juniper.net/support/eol/junos.html>

---

## Downloading Software Upgrades from Juniper Networks

Follow these steps to download software upgrades from Juniper Networks:

1. Using a Web browser, follow the links to the download URL on the Juniper Networks Web page. Depending on your location, select either **Canada and U.S. Version** or **Worldwide Version**:
  - <https://www.juniper.net/support/csc/swdist-domestic/> (customers in the United States and Canada)
  - <https://www.juniper.net/support/csc/swdist-ww/> (all other customers)
2. Log in to the Juniper Networks Web site using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Using the J-Web interface or the CLI, select the appropriate junos-j-series software package for your application. For information about JUNOS software packages, see “Upgrade and Downgrade Overview” on page 85.
4. Download the software to a local host or to an internal software distribution site.



**NOTE:** For downloads to J-series Services Routers with a 256-MB compact flash, see “Special Instructions for J-series Routers with a 256-MB Compact Flash” on page 92.

---

## Installing Software Upgrades with the J-Web Interface

If your router has at least a 512-MB compact flash, you can use the J-Web interface to install software upgrades from a remote server using FTP or HTTP, or by uploading the software image to the router. This section contains the following topics:

- Installing Software Upgrades from a Remote Server on page 87
- Installing Software Upgrades by Uploading Files on page 88

### Installing Software Upgrades from a Remote Server

If your router has at least a 512-MB compact flash, you can use the J-Web interface to install software packages on the router that are retrieved with FTP or HTTP from the location specified.

To install software upgrades from a remote server:

1. Download the software package as described in “Downloading Software Upgrades from Juniper Networks” on page 87.
2. In the J-Web interface, select **Manage > Software > Install Package**.
3. On the Install Package page, enter information into the fields described in Table 2 on page 88.
4. Click **Fetch and Install Package**. The software is activated after the router has rebooted.

**Table 2: Install Package Summary**

Field	Function	Your Action
Package Location (required)	Specifies the FTP or HTTP server, file path, and software package name.	Type the full address of the software package location on the FTP or HTTP server—one of the following:  <code>ftp://hostname/pathname/package-name</code> <code>http://hostname/pathname/package-name</code>
User	Specifies the username, if the server requires one.	Type the username.
Password	Specifies the password, if the server requires one.	Type the password.
Reboot If Required	If this box is checked, the router is automatically rebooted when the upgrade is complete.	Check the box if you want the router to reboot automatically when the upgrade is complete.

### Installing Software Upgrades by Uploading Files

If your router has at least a 512-MB compact flash, you can use the J-Web interface to install software packages uploaded from your computer to the router.

To install software upgrades by uploading files:

1. Download the software package as described in “Downloading Software Upgrades from Juniper Networks” on page 87.
2. In the J-Web interface, select **Manage > Software > Upload Package**.
3. On the Upload Package page, enter information into the fields described in Table 3 on page 88.
4. Click **Upload Package**. The software is activated after the router has rebooted.

**Table 3: Upload Package Summary**

Field	Function	Your Action
File to Upload (required)	Specifies the location of the software package.	Type the location of the software package, or click <b>Browse</b> to navigate to the location.
Reboot If Required	If this box is checked the router is automatically rebooted when the upgrade is complete.	Select the check box if you want the router to reboot automatically when the upgrade is complete.

### Installing Software Upgrades with the CLI

You can use the CLI to install software upgrades from a remote server using FTP or by downloading the software image to the router. If your router has a 256-MB compact

flash, see “Special Instructions for J-series Routers with a 256-MB Compact Flash” on page 92.

This section contains the following topics:

- Installing Software Upgrades by Downloading Files on page 89
- Installing Software Upgrades from a Remote Server on page 90

## Installing Software Upgrades by Downloading Files

To install software upgrades by downloading files to the router:

1. Download the JUNOS software package to the router using the following command:

```
user@host> file copy source destination
```

Replace *source* with one of the following paths:

- `ftp://hostname/pathname/package-name`
- or
- `http://hostname/pathname/package-name`

Replace *destination* with the path to the destination directory on the router. We recommend the `/var/tmp` directory.

If you had configured the unused swap partition using the `upgrade-helper` script (as described in “Configuring the Unused Swap Partition” on page 94), make sure to copy the software package to the `/var/tmp/upgrade` directory.

2. Install the new package on the Services Router, entering the following command in operational mode in the CLI:

```
user@host> request system software add validate unlink no-copy source
```

Replace *source* with `/pathname/package-name` (for example, `/var/tmp/junos-jsr-8.5R2.1.tar.gz`).

By default, the `request system software add` command uses the `validate` option to validate the software package against the current configuration as a prerequisite to adding the software package. This validation ensures that the router can reboot successfully after the software package is installed. This is the default behavior when you are adding a software package.

The `unlink` option removes the package at the earliest opportunity so that the router has enough room to complete the installation.

(Optional) The `no-copy` option specifies that a software package is installed, but a copy of the package is not saved. Include this option if you do not have enough space on the compact flash to perform an upgrade that keeps a copy of the package on the router.

3. After the software package is installed, reboot the router:

```
user@host> request system reboot
```

When the reboot is complete, the router displays the login prompt.

## Installing Software Upgrades from a Remote Server

To install the software upgrades from a remote server:

1. Install the JUNOS software package on the Services Router, entering the following command in operational mode in the CLI:

```
user@host> request system software add validate unlink no-copy source
```

Replace *source* with one of the following paths:

- `ftp://hostname/pathname/package-name`
- or
- `http://hostname/pathname/package-name`

By default, the `request system software add` command uses the `validate` option to validate the software package against the current configuration as a prerequisite to adding the software package. This validation ensures that the router can reboot successfully after the software package is installed. This is the default behavior when you are adding a software package.

The `unlink` option removes the package at the earliest opportunity so that the router has enough room to complete the installation.

(Optional) The `no-copy` option specifies that a software package is installed, but a copy of the package is not saved. Include this option if you do not have enough space on the compact flash to perform an upgrade that keeps a copy of the package on the router.

2. After the software package is installed, reboot the router:

```
user@host> request system reboot
```

When the reboot is complete, the router displays the login prompt.

## Downgrade Instructions

This section contains the following topics:

- Downgrading the Software with the J-Web Interface on page 91
- Downgrading the Software with the CLI on page 91



**NOTE:** Juniper Networks supports direct software downgrades for a maximum of three releases. For example, if your routing platform is running JUNOS Release 7.6, you can typically downgrade without problems to Release 7.3. If you attempt to downgrade more than three releases and validation of your configuration fails, we recommend downgrading to an intermediate release first before downgrading to the desired release.

---

### Downgrading the Software with the J-Web Interface

You can downgrade the software from the J-Web interface. For the changes to take effect, you must reboot the router.

To downgrade software:

1. In the J-Web interface, select **Manage > Software > Downgrade**. The image of the previous software version (if any) is displayed on this page.



**NOTE:** After you perform this operation, you cannot undo it.

2. Select **Downgrade** to downgrade to the previous version of the software or **Cancel** to cancel the downgrade process.
3. When the downgrade process is complete, for the new software to take effect, select **Manage > Reboot** from the J-Web interface to reboot the router.

After you downgrade the software, the previous release is loaded, and you cannot reload the running version of software again. To downgrade to an earlier version of software, follow the procedure for upgrading, using the JUNOS software image labeled with the appropriate release.

### Downgrading the Software with the CLI

You can revert to the previous version of software using the `request system software rollback` command in the CLI. For the changes to take effect, you must reboot the router. To downgrade to an earlier version of software, follow the procedure for upgrading, using the JUNOS software image labeled with the appropriate release.

To downgrade software with the CLI:

1. Enter the `request system software rollback` command to return to the previous JUNOS software version:

```
user@host> request system software rollback
```

The previous software version is now ready to become active when you next reboot the router.

2. Reboot the router:

```
user@host> request system reboot
```

The router is now running the previous version of the software. To downgrade to an earlier version of software, follow the procedure for upgrading, using the JUNOS software image labeled with the appropriate release.

### **Special Instructions for J-series Routers with a 256-MB Compact Flash**

J-series Services Routers with a 256-MB compact flash might need more flash memory space for an upgrade.

To provide enough space for an upgrade:

- Clean up files on the router (see “Cleaning Up Files” on page 92).
- Verify the available compact flash space (see “Verifying Available Compact Flash Space” on page 93).
- If required, increase the compact flash space, (see “Increasing the Compact Flash Space” on page 93).

### **Cleaning Up Files**

To clean up files, you use CLI commands to delete the backup software image, rotate log files, and remove other unnecessary files.

When you upgrade software on the router, it creates a backup image of the software that was previously installed. To create enough space on a 256-MB compact flash for an upgrade, use the `request system software delete backup` command to delete this image. In addition, use the `request system storage cleanup` command to rotate log files and delete unnecessary files.



**NOTE:** To review the list of files that can be deleted without actually deleting files, you can use the `request system storage cleanup dry-run` command.

---

To delete the backup software image, rotate log files, and delete unneeded files:

1. From operational mode in the CLI, enter the following command:

```
user@host> request system software delete backup
```

2. Enter **yes** when prompted:

```
Delete backup system software package [yes,no] (no) yes
```

3. Enter the following command:

```
user@host> request system storage cleanup
```

The router rotates log files and displays the files that you can delete.

4. Enter **yes** at the prompt to delete the files.
5. Delete any files that you created by entering the following command:

```
user@host> file delete filename
```

Replace *filename* with the name of the file or directory to delete.

- Verify that you have enough space on the compact flash to successfully upgrade (see “Verifying Available Compact Flash Space” on page 93).

## Verifying Available Compact Flash Space

Before you start the upgrade, verify that you have enough space on the compact flash to successfully upgrade.

To see how much space is available on the compact flash, use the CLI operational mode command `show system storage`:

```
user@host show system storage
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     213M     119M      92M    57%      /
devfs           1.0K     1.0K       0B    100%    /dev
devfs           1.0K     1.0K       0B    100%    /dev/
/dev/md0        155M     155M       0B    100%    /junos
/cf            213M     119M      92M    57%    /junos/cf
devfs           1.0K     1.0K       0B    100%    /junos/dev/
procfs         4.0K     4.0K       0B    100%    /proc
/dev/bo0s1e     24M      16K       24M     0%    /config
/dev/md1        168M     7.2M     147M     5%    /mfs
/dev/md2         58M      42K       53M     0%    /jail/tmp
/dev/md3         7.7M    100K       7.0M     1%    /jail/var/etc
devfs           1.0K     1.0K       0B    100%    /jail/dev
/dev/md4         1.9M     6.0K      1.7M     0%    /jail/html/oem
```

The `show system storage` command output displays information about the root file system on the compact flash on the line that contains only a forward slash (/) in the **Mounted on** column. In this example, the compact flash has 92 MB of available space.

If the `show system storage` command output displays:

- Available compact flash space—135 MB or more. See “Installing Software Upgrades with the CLI” on page 88 to proceed with the upgrade.
- Available compact flash space—less than 135 MB. See “Increasing the Compact Flash Space” on page 93 to increase the compact flash space.

## Increasing the Compact Flash Space



**NOTE:** On J-series Services Routers running JUNOS Release 8.2 or later, you can no longer specify the internal compact flash as the medium used to store system software failure memory snapshots when using the `set system dump-device` CLI command. For J4350 or J6350 Services Routers, you need to specify a USB storage device (`usb` option) as the medium. For J2320 and J2350 Services Routers, you can specify a USB storage device (`usb` option) or the external compact flash (`removable-compact-flash` option) as the medium.

To increase the compact flash space:

- If you have physical access to the router, remove the swap partition (see “Removing the Swap Partition” on page 94).
- If you do not have physical access to the router, download the **upgrade-helper** script to configure the unused swap partition (see “Configuring the Unused Swap Partition” on page 94).

## Removing the Swap Partition

To remove the swap partition:

1. Insert a Juniper Networks-supported 256-MB USB storage device into an available USB port of the Services Router to be upgraded.
2. From operational mode in the CLI, enter the following command:

```
user@host> request system snapshot as-primary partition swap-size 0 media usb
```

3. Enter the following command:

```
user@host> request system reboot media usb
```

This command reboots the router and boots from the USB storage device with the original configuration file intact. After rebooting, the router is online and uses the configuration file as the running configuration.

4. Enter the following command:

```
user@host> request system snapshot as-primary partition swap-size 0 media compact-flash
```

This command repartitions the internal compact flash so that it has no swap partition.

5. Enter the following command:

```
user@host> request system reboot media compact-flash
```

This command reboots the router from the internal compact flash. After rebooting, the router is online with your running configuration, but the swap partition on the compact flash is removed.

6. Remove the USB storage device.
7. See “Installing Software Upgrades with the CLI” on page 88 to proceed with the upgrade.

## Configuring the Unused Swap Partition

To configure the unused swap partition:

1. In your Web browser, type the following URL. When prompted, use the username and password supplied to you by Juniper Networks representatives to download the `upgrade-helper` script to your local server.

```
https://download.juniper.net/software/junos/specials/upgrade-helper.gz
```

2. Start a UNIX-level shell and log in as a root user.
3. Enter the CLI and from the operational mode copy the `upgrade-helper` script to the `root` directory on your router:

```
user@host> file copy source destination
```

Replace *source* with the path to the script on your local server.

Replace *destination* with the destination directory: `/root`.

4. Exit the CLI environment and create a UNIX-level shell:

```
user@host> start shell
```

5. Use the compression utility `gunzip` to decompress the downloaded script. The `gunzip` utility is available on your router in the `/usr/bin/gunzip` directory.
6. Execute the script:

```
root@host% sh ./upgrade-helper
Upgrade helper script started
ATTENTION: PLEASE RUN THIS SCRIPT AGAIN IMMEDIATELY AFTER REBOOTING.
Rebooting system.
```

The system reboots (in no more than 10 seconds) without a swap partition.

7. Execute the `upgrade-helper` script again immediately after rebooting.
8. See “Installing Software Upgrades by Downloading Files” on page 89 to proceed with the upgrade.

## List of Technical Publications

---

Table 4 on page 96 lists the software and hardware guides and release notes for Juniper Networks J-series, M-series, MX-series, and T-series routing platforms and describes the contents of each document. Table 5 on page 100 lists the books included in the *Network Operations Guide* series. Table 6 on page 100 lists the manuals and release notes supporting JUNOS software with enhanced services. All documents are available at <http://www.juniper.net/techpubs/>.

Table 7 on page 102 lists additional books on Juniper Networks solutions that you can order through your bookstore. A complete list of such books is available at <http://www.juniper.net/books>.

**Table 4: Technical Documentation for Supported Routing Platforms**

Book	Description
<b>JUNOS Software for Supported Routing Platforms</b>	
<i>Access Privilege</i>	Explains how to configure access privileges in user classes by using permission flags and regular expressions. Lists the permission flags along with their associated command-line interface (CLI) operational mode commands and configuration statements.
<i>Class of Service</i>	Provides an overview of the class-of-service (CoS) functions of the JUNOS software and describes how to configure CoS features, including configuring multiple forwarding classes for transmitting packets, defining which packets are placed into each output queue, scheduling the transmission service level for each queue, and managing congestion through the random early detection (RED) algorithm.
<i>CLI User Guide</i>	Describes how to use the JUNOS command-line interface (CLI) to configure, monitor, and manage Juniper Networks routing platforms. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .
<i>Feature Guide</i>	Provides a detailed explanation and configuration examples for several of the most complex features in the JUNOS software.
<i>High Availability</i>	Provides an overview of hardware and software resources that ensure a high level of continuous routing platform operation and describes how to configure high availability (HA) features such as nonstop active routing (NSR) and graceful Routing Engine switchover (GRES).
<i>MPLS Applications</i>	Provides an overview of traffic engineering concepts and describes how to configure traffic engineering protocols.
<i>Multicast Protocols</i>	Provides an overview of multicast concepts and describes how to configure multicast routing protocols.
<i>Multiplay Solutions</i>	Describes how you can deploy IPTV and voice over IP (VoIP) services in your network.
<i>MX-series Solutions Guide</i>	Describes common configuration scenarios for the Layer 2 features supported on the MX-series routers, including basic bridged VLANs with normalized VLAN tags, aggregated Ethernet links, bridge domains, Multiple Spanning Tree Protocol (MSTP), and integrated routing and bridging (IRB).
<i>Network Interfaces</i>	Provides an overview of the network interface functions of the JUNOS software and describes how to configure the network interfaces on the routing platform.
<i>Network Management</i>	Provides an overview of network management concepts and describes how to configure various network management features, such as SNMP and accounting options.
<i>Policy Framework</i>	Provides an overview of policy concepts and describes how to configure routing policy, firewall filters, and forwarding options.

**Table 4: Technical Documentation for Supported Routing Platforms** (continued)

Book	Description
<i>Protected System Domain</i>	Provides an overview of the JCS 1200 platform and the concept of Protected System Domains (PSDs). The JCS 1200 platform, which contains up to six redundant pairs of Routing Engines running JUNOS software, is connected to a T320 router or to a T640 or T1600 routing node. To configure a PSD, you assign any number of Flexible PIC concentrators (FPCs) in the T-series routing platform to a pair of Routing Engines on the JCS 1200 platform. Each PSD has the same capabilities and functionality as a physical router, with its own control plane, forwarding plane, and administration.
<i>Routing Protocols</i>	Provides an overview of routing concepts and describes how to configure routing, routing instances, and unicast routing protocols.
<i>Secure Configuration Guide for Common Criteria and JUNOS-FIPS</i>	Provides an overview of secure Common Criteria and JUNOS-FIPS protocols for the JUNOS software and describes how to install and configure secure Common Criteria and JUNOS-FIPS on a routing platform.
<i>Services Interfaces</i>	Provides an overview of the services interfaces functions of the JUNOS software and describes how to configure the services interfaces on the router.
<i>Software Installation and Upgrade Guide</i>	Describes the JUNOS software components and packaging and explains how to initially configure, reinstall, and upgrade the JUNOS system software. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .
<i>Subscriber Access</i>	Provides an overview of the subscriber access features of the JUNOS software and describes how to configure subscriber access support on the router, including dynamic profiles, class of service, AAA, and access methods.
<i>System Basics</i>	Describes Juniper Networks routing platforms and explains how to configure basic system parameters, supported protocols and software processes, authentication, and a variety of utilities for managing your router on the network.
<i>VPNs</i>	Provides an overview and describes how to configure Layer 2 and Layer 3 virtual private networks (VPNs), virtual private LAN service (VPLS), and Layer 2 circuits. Provides configuration examples.
<b>JUNOS References</b>	
<i>Hierarchy and RFC Reference</i>	Describes the JUNOS configuration mode commands. Provides a hierarchy reference that displays each level of a configuration hierarchy, and includes all possible configuration statements that can be used at that level. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .
<i>Interfaces Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot interfaces.
<i>Routing Protocols and Policies Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot routing policies and protocols, including firewall filters.

**Table 4: Technical Documentation for Supported Routing Platforms (continued)**

Book	Description
<i>System Basics and Services Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot system basics, including commands for real-time monitoring and route (or path) tracing, system software management, and chassis management. Also describes commands for monitoring and troubleshooting services such as class of service (CoS), IP Security (IPsec), stateful firewalls, flow collection, and flow monitoring.
<i>System Log Messages Reference</i>	Describes how to access and interpret system log messages generated by JUNOS software modules and provides a reference page for each message.
<b>J-Web User Guide</b>	
<i>J-Web Interface User Guide</i>	Describes how to use the J-Web graphical user interface (GUI) to configure, monitor, and manage Juniper Networks routing platforms.
<b>JUNOS API and Scripting Documentation</b>	
<i>JUNOScript API Guide</i>	Describes how to use the JUNOScript application programming interface (API) to monitor and configure Juniper Networks routing platforms.
<i>JUNOS XML API Configuration Reference</i>	Provides reference pages for the configuration tag elements in the JUNOS XML API.
<i>JUNOS XML API Operational Reference</i>	Provides reference pages for the operational tag elements in the JUNOS XML API.
<i>NETCONF API Guide</i>	Describes how to use the NETCONF API to monitor and configure Juniper Networks routing platforms.
<i>JUNOS Configuration and Diagnostic Automation Guide</i>	Describes how to use the commit script and self-diagnosis features of the JUNOS software. This guide explains how to enforce custom configuration rules defined in scripts, how to use commit script macros to provide simplified aliases for frequently used configuration statements, and how to configure diagnostic event policies.
<b>Hardware Documentation</b>	
<i>Hardware Guide</i>	Describes how to install, maintain, and troubleshoot routing platforms and components. Each platform has its own hardware guide.
<i>PIC Guide</i>	Describes the routing platform's Physical Interface Cards (PICs). Each platform has its own PIC guide.
<i>DPC Guide</i>	Describes the Dense Port Concentrators (DPCs) for all MX-series routers.
<b>JUNOScope Documentation</b>	
<i>JUNOScope Software User Guide</i>	Describes the JUNOScope software graphical user interface (GUI), how to install and administer the software, and how to use the software to manage routing platform configuration files and monitor routing platform operations.

**Table 4: Technical Documentation for Supported Routing Platforms (continued)**

Book	Description
<b>Advanced Insight Solutions (AIS) Documentation</b>	
<i>Advanced Insight Solutions Guide</i>	Describes the Advanced Insight Manager (AIM) application, which provides a gateway between JUNOS devices and Juniper Support Systems (JSS) for case management and intelligence updates. Explains how to run AI-Scripts on Juniper Networks devices.
<b>J-series Routing Platform Documentation</b>	
<i>Getting Started Guide</i>	Provides an overview, basic instructions, and specifications for J-series routing platforms. The guide explains how to prepare your site for installation, unpack and install the router and its components, install licenses, and establish basic connectivity. Use the <i>Getting Started Guide</i> for your router model.
<i>Basic LAN and WAN Access Configuration Guide</i>	Explains how to configure the interfaces on J-series Services Routers for basic IP routing with standard routing protocols, ISDN backup, and digital subscriber line (DSL) connections.
<i>Advanced WAN Access Configuration Guide</i>	Explains how to configure J-series Services Routers in virtual private networks (VPNs) and multicast networks, configure data link switching (DLSw) services, and apply routing techniques such as policies, stateless and stateful firewall filters, IP Security (IPsec) tunnels, and class-of-service (CoS) classification for safer, more efficient routing.
<i>Administration Guide</i>	Shows how to manage users and operations, monitor network performance, upgrade software, and diagnose common problems on J-series Services Routers.
<b>Release Notes</b>	
<i>JUNOS Release Notes</i>	Summarize new features and known problems for a particular software release, provide corrections and updates to published JUNOS, JUNOScript, and NETCONF manuals, provide information that might have been omitted from the manuals, and describe upgrade and downgrade procedures.
<i>Hardware Release Notes</i>	Describe the available documentation for the routing platform and summarize known problems with the hardware and accompanying software. Each platform has its own release notes.
<i>JUNOScope Release Notes</i>	Contain corrections and updates to the published JUNOScope manual, provide information that might have been omitted from the manual, and describe upgrade and downgrade procedures.
<i>AIS Release Notes</i>	Summarize AIS new features and guidelines, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide initial setup, upgrade, and downgrade procedures.
<i>AIS AI-Scripts Release Notes</i>	Summarize AI-Scripts new features, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide instructions for automatic and manual installation, including deleting and rolling back.

**Table 4: Technical Documentation for Supported Routing Platforms** (continued)

Book	Description
<i>J-series Services Router Release Notes</i>	Briefly describe Services Router features, identify known hardware problems, and provide upgrade and downgrade instructions.

**Table 5: JUNOS Software Network Operations Guides**

Book	Description
<i>Baseline</i>	Describes the most basic tasks for running a network using Juniper Networks products. Tasks include upgrading and reinstalling JUNOS software, gathering basic system management information, verifying your network topology, and searching log messages.
<i>Interfaces</i>	Describes tasks for monitoring interfaces. Tasks include using loopback testing and locating alarms.
<i>MPLS</i>	Describes tasks for configuring, monitoring, and troubleshooting an example MPLS network. Tasks include verifying the correct configuration of the MPLS and RSVP protocols, displaying the status and statistics of MPLS running on all routing platforms in the network, and using the layered MPLS troubleshooting model to investigate problems with an MPLS network.
<i>MPLS Log Reference</i>	Describes MPLS status and error messages that appear in the output of the <code>show mpls lsp extensive</code> command. The guide also describes how and when to configure Constrained Shortest Path First (CSPF) and RSVP trace options, and how to examine a CSPF or RSVP failure in a sample network.
<i>MPLS Fast Reroute</i>	Describes operational information helpful in monitoring and troubleshooting an MPLS network configured with fast reroute (FRR) and load balancing.
<i>Hardware</i>	Describes tasks for monitoring M-series and T-series routing platforms.

To configure and operate a J-series Services Router or SRX-series services gateway running JUNOS software with enhanced services, you must also use the configuration statements and operational mode commands documented in JUNOS configuration guides and command references. To configure and operate a WX Integrated Services Module, you must also use WX documentation.

**Table 6: JUNOS Software with Enhanced Services Documentation**

Book	Description
<b>All Platforms</b>	
<i>JUNOS Software Interfaces and Routing Configuration Guide</i>	Explains how to configure J-series and SRX-series interfaces for basic IP routing with standard routing protocols, ISDN service, firewall filters (access control lists), and class-of-service (CoS) traffic classification.

**Table 6: JUNOS Software with Enhanced Services Documentation (continued)**

Book	Description
<i>JUNOS Software Security Configuration Guide</i>	Explains how to configure and manage security services such as stateful firewall policies, IP Security (IPsec) virtual private networks (VPNs), firewall screens, Network Address Translation (NAT), Public Key Cryptography, and Application Layer Gateways (ALGs).
<i>JUNOS Software Administration Guide</i>	Shows how to monitor J-series and SRX-series devices and routing operations, firewall and security services, system alarms and events, and network performance. This guide also shows how to administer user authentication and access, upgrade software, and diagnose common problems.
<i>JUNOS Software CLI Reference</i>	Provides the complete JUNOS software with enhanced services configuration hierarchy and describes the configuration statements and operational mode commands not documented in the standard JUNOS manuals.
<b>J-series Only</b>	
<i>JUNOS Software with Enhanced Services Design and Implementation Guide</i>	Provides guidelines and examples for designing and implementing IPsec VPNs, firewalls, and routing on J-series Services Routers running JUNOS software with enhanced services.
<i>JUNOS Software with Enhanced Services Quick Start</i>	Explains how to quickly set up a J-series Services Router. This document contains router declarations of conformity.
<i>JUNOS Software with Enhanced Services J-series Services Router Hardware Guide</i>	Provides an overview, basic instructions, and specifications for J-series Services Routers. This guide explains how to prepare a site, unpack and install the router, replace router hardware, and establish basic router connectivity. This guide contains hardware descriptions and specifications.
<i>JUNOS Software with Enhanced Services Migration Guide</i>	Provides instructions for migrating an SSG device running ScreenOS software or a J-series Services Router running the JUNOS software to JUNOS software with enhanced services.
<i>WXC Integrated Services Module Installation and Configuration Guide</i>	Explains how to install and initially configure a WXC Integrated Services Module in a J-series Services Router for application acceleration.
<i>JUNOS Software with Enhanced Services Release Notes</i>	Summarizes new features and known problems for a particular release of JUNOS software with enhanced services on J-series Services Routers, including J-Web interface features and problems. The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions for JUNOS software with enhanced services.
<b>SRX-series Only</b>	

**Table 6: JUNOS Software with Enhanced Services Documentation (continued)**

Book	Description
<i>JUNOS Software for SRX-series Services Gateway Release Notes</i>	Summarizes new features and known problems for a particular release of JUNOS software with enhanced services on SRX-series services gateways, including J-Web interface features and problems. The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions for JUNOS software with enhanced services.

**Table 7: Additional Books Available Through <http://www.juniper.net/books>**

Book	Description
<i>Interdomain Multicast Routing</i>	Provides background and in-depth analysis of multicast routing using Protocol Independent Multicast sparse mode (PIM SM) and Multicast Source Discovery Protocol (MSDP); details any-source and source-specific multicast delivery models; explores multiprotocol BGP (MBGP) and multicast IS-IS; explains Internet Gateway Management Protocol (IGMP) versions 1, 2, and 3; lists packet formats for IGMP, PIM, and MSDP; and provides a complete glossary of multicast terms.
<i>JUNOS Cookbook</i>	Provides detailed examples of common JUNOS software configuration tasks, such as basic router configuration and file management, security and access control, logging, routing policy, firewalls, routing protocols, MPLS, and VPNs.
<i>MPLS-Enabled Applications</i>	Provides an overview of Multiprotocol Label Switching (MPLS) applications (such as Layer 3 virtual private networks [VPNs], Layer 2 VPNs, virtual private LAN service [VPLS], and pseudowires), explains how to apply MPLS, examines the scaling requirements of equipment at different points in the network, and covers the following topics: point-to-multipoint label switched paths (LSPs), DiffServ-aware traffic engineering, class of service, interdomain traffic engineering, path computation, route target filtering, multicast support for Layer 3 VPNs, and management and troubleshooting of MPLS networks.
<i>OSPF and IS-IS: Choosing an IGP for Large-Scale Networks</i>	Explores the full range of characteristics and capabilities for the two major link-state routing protocols: Open Shortest Path First (OSPF) and IS-IS. Explains architecture, packet types, and addressing; demonstrates how to improve scalability; shows how to design large-scale networks for maximum security and reliability; details protocol extensions for MPLS-based traffic engineering, IPv6, and multitopology routing; and covers troubleshooting for OSPF and IS-IS networks.
<i>Routing Policy and Protocols for Multivendor IP Networks</i>	Provides a brief history of the Internet, explains IP addressing and routing (Routing Information Protocol [RIP], OSPF, IS-IS, and Border Gateway Protocol [BGP]), explores ISP peering and routing policies, and displays configurations for both Juniper Networks and other vendors' routers.
<i>The Complete IS-IS Protocol</i>	Provides the insight and practical solutions necessary to understand the IS-IS protocol and how it works by using a multivendor, real-world approach.

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net), or fill out the documentation feedback form at

<http://www.juniper.net/techpubs/docbug/docbugreport.html>. If you are using e-mail, be sure to include the following information with your comments:

- Document name
- Document part number
- Page number
- Software release version

## 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 postsales 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 us at <http://www.juniper.net/support/requesting-support.html>.

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the `gzip` utility, rename the file to include your company name, and copy it to `ftp.juniper.net:pub/incoming`. Then send the filename, along with software version information (the output of the `show version` command) and the configuration, to `support@juniper.net`. For documentation issues, fill out the bug report form located at <http://www.juniper.net/techpubs/docbug/docbugreport.html>.

## Revision History

---

18 August 2008—Revision 6, JUNOS 8.5R4

25 April 2008—Revision 5, JUNOS 8.5R3

25 March 2008—Revision 4, JUNOS 8.5R2

12 February 2008—Revision 3, JUNOS 8.5R2.

21 November 2007—Revision 2, JUNOS 8.5R1.

16 November 2007—Revision 1, JUNOS 8.5R1.

Copyright © 2008, Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, JUNOS, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark 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 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.