JUNOS® 9.3 Software Release Notes

Release 9.3R4
3 December 2010
Revision 6

These release notes accompany Release 9.3R4 of the JUNOS software. They describe device documentation and known problems with the software. JUNOS software runs on all Juniper Networks M-series, MX-series, and T-series routing platforms, SRX-series Services Gateways, J-series Services Routers, and EX-series switches.

You can also find these release notes on the Juniper Networks JUNOS Software Documentation Web page, which is located at http://www.juniper.net/techpubs/software/junos/.

Contents

JUNOS Software Release Notes for M-series, MX-series, and T-series Routing Platforms

Features in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms

The following features have been added to JUNOS Release 9.3. Following the description is the title of the manual or manuals to consult for further information. For a complete list of manuals, see Table 11 on page 160, Table 12 on page 164, and Table 14 on page 166.

Hardware

- **Extended FPC support for MultiServices 500 PICs**—MultiServices 500 (Type 3) PICs are now supported on T640-FPC3-ES Flexible PIC Concentrators (FPCs), which can be installed in T640 and T1600 routing platforms. All the existing functionality, including both Layer 2 and Layer 3 services packages, is supported. There are no changes to the command-line interface (CLI) or the behavior of the PIC. [*T640 PIC Guide*, *T1600 PIC Guide*]
- **New Enhanced IQ PICs with SFP (M40e, M120, and T-series routing platforms)**—The following Enhanced Intelligent Queuing (IQ) PICs with small form-factor pluggable transceivers (SFP) are available:
  - 4-port Channelized DS3/E3 Enhanced IQ PIC
  - 2-port Channelized SONET/SDH OC3/STM1 Enhanced IQ PIC
  - 1-port Channelized SONET/SDH OC12/STM4 Enhanced IQ PIC
  - 4-port DS3/E3 Enhanced IQ PIC
  - 4-port SONET/SDH OC3/STM1 Enhanced IQ PIC
  - 1-port SONET/SDH OC12/STM4 Enhanced IQ PIC

  The Enhanced IQ PICs support all the same features as existing IQ PICs. The valid configuration statements are also the same; for some options, limits and ranges of values are different to support augmented capabilities. [*PIC Guides*, *CoS*, *Network Interfaces*]

- **New Circuit Emulation PICs (M7i, M10i, and M40e routers)**—New Circuit Emulation PICs support structure-agnostic time-division multiplexing (TDM) over packet (SAToP) as defined in RFC 4553. The PICs provide Layer 2 circuit emulation service, which is used mainly in mobile backhaul applications. The following PICs are available:
  - 12-port T1/E1 Circuit Emulation PIC
  - 4-port SONET/SDH OC3/STM1 Circuit Emulation PIC

  You can configure the 12-port T1/E1 PIC for either 12 T1-mode or 12 E1-mode interfaces, but not a combination of both interface types. To specify the mode for all interfaces on the PIC, include either the framing t1 or framing e1 statement at the [edit chassis fpc *slot-number* pic *slot-number*] hierarchy level. To configure each of the 12 interfaces, include either the interface-type t1 statement at the [edit interfaces ct1-*fpc/pic/port* no-partition] hierarchy level (for T1 mode) or the interface-type ce1 statement at the [edit interfaces ce1-*fpc/pic/port* no-partition] hierarchy level (for E1 mode). [*PIC Guides*, *Network Interfaces*, *System Basics*]
- **New MultiServices DPC (MX-series Ethernet services routers)**—The MultiServices Dense Port Concentrator (DPC) supports the following Layer 3 services:
  - Active flow monitoring
  - Generic routing encapsulation (GRE) tunnels (including GRE key and fragmentation)
  - Intrusion detection service (IDS)
  - IP Security (IPsec)
  - Network address translation (NAT)
  - Real-time performance monitoring (RPM)
  - Stateful firewall

  All functionality available for these features on Adaptive Services and MultiServices PICs is also available on MultiServices DPCs and is configured by including the existing statements at the [edit services] hierarchy level. All existing commands for monitoring the features are also supported. The show chassis commands represent MultiServices DPCs as MS-DPC or MS-DPC PIC. [*DPC Guide*, *Network Interfaces*, *Services Interfaces*]

Software Installation and Upgrade

- **License support for CoA**—The RADIUS Change-of-Authorization (CoA) feature is part of the Subscriber Access Feature Pack License. As with other licensed JUNOS features, you can commit a configuration that includes the CoA feature without having a license and use the feature for a 30-day trial period (dated from first actual use of the feature). To display license key information, issue the show system license command. [*Software Installation and Upgrade*]

User Interface and Configuration

- **"Logical system" replaces "logical router"**—Beginning in JUNOS Release 9.3, the term *logical system* and the JUNOS name logical-system replace their previous counterparts *logical router* and logical-router in all contexts, including configuration statements, command names and output fields, error and log messages, and SNMP MIB object names. The [edit logical-routers] hierarchy level is now [edit logical-systems].

  The name logical-router is still accepted in JUNOS Release 9.3, to enable continued use of scripts and other tools that use the previous name. Similarly, in JUNOS Releases 8.3R2 through 9.2, logical-system is accepted as an alternate to logical-router. We encourage you to update your JUNOS environment to use the new term. [All JUNOS documentation]

Interfaces and Chassis

- **External clock synchronization (M120 routers)**—Enables you to synchronize an M120 router's internal Stratum 3 clock module to a standard external timing source such as Building Integrated Timing System (BITS), SDH Equipment Timing Source (SETS), or an equivalent quality timing source. External clock synchronization is supported for T1, E1, and SONET/SDH interfaces. To configure external clock synchronization, include the synchronization statement at the [edit chassis] hierarchy level. To view information about the external clock synchronization, issue the show chassis synchronization command. To change the external clock synchronization time source, issue the request chassis synchronization switch command. [*System Basics*, *System Basics Command Reference*]
- **Interfaces shared by PSDs**—Enables multiple Protected System Domains (PSDs) to receive and forward traffic on the interfaces on a PIC. (A PSD is a Routing Engine [or pair of redundant Routing Engines] on the JCS 1200 platform that is associated with one or more FPCs on a T-series routing platform.) Any FPC that is not assigned to a specific PSD can be used to host shared interfaces. In JUNOS Release 9.3, interfaces on SONET/SDH PICs only can be shared.

  On the Root System Domain (RSD), you configure multiple logical interfaces on a physical SONET/SDH interface (the *shared uplink* interface) and assign each logical interface to a specific PSD. On the PSD, you include a corresponding configuration stanza for the physical interface and the logical interfaces, peering each logical interface with a logical interface on a Tunnel PIC managed by the PSD. The Tunnel PIC transports packets between that logical interface on the PSD and the physical interface on the RSD. In the configuration on both the RSD

and PSD, you must specify Frame Relay encapsulation for the physical interface (but point-to-multipoint Frame Relay encapsulation is not supported).

On the RSD, assign each logical interface to a PSD by including the new uplink-shared-with psd*n* statement at the [edit interfaces so-*fpc/pic/port*.*logical-unit-number*] hierarchy level for the logical interface.

On the PSD, include the following new statements at the indicated hierarchy levels:

- The shared-uplink statement at the [edit interfaces so-*fpc/pic/port*] hierarchy level, to identify the physical SONET interface acting as the shared uplink
- The ut-*fpc/pic/port* statement at the [edit interfaces] hierarchy level, to configure the physical interface on the Tunnel PIC as an uplink tunnel interface

You must also include other existing statements to complete the interface configuration on the RSD and PSDs; for details, see the PSD documentation.

When applied to shared interfaces, JUNOS features that are configured on logical interfaces—such as CoS classifiers and rewrite rules, firewall filters, and policers—are configured on the PSD. With JUNOS Release 9.3, only output filtering is supported. JUNOS features that are configured on physical interfaces, such as drop profiles and scheduler maps, are configured on the RSD.

The following new fields in the output from the show interfaces so-*fpc/pic/port* command display information about shared interfaces:

- Shared-uplink—Located in the Physical interface: section; indicates whether the routing domain owns the shared interface.
- Shared uplink—Located in the Logical interface: section; includes these fields:
  - shared with—(RSD only) Names the PSD that owns the logical shared interface. For example, psd3.
  - peer interface—(PSD only) Names the logical tunnel interface that peers with the logical shared interface. For example, ut-2/1/0.2.
  - tunnel token—Specifies the receive (Rx) and transmit (Tx) tunnel tokens. For example, Rx: 5.519, Tx: 13.514.

[*Protected System Domain*]

- **Finer control of active and standby LACP links (MX-series routers)**—Enables you to better control the behavior of active and standby links on aggregated Ethernet interfaces by more precisely defining Link Aggregation Control Protocol (LACP) system and port priorities. To configure LACP system priority for all aggregated Ethernet interfaces on a router, include the new system-priority statement at the [edit chassis aggregated-devices ethernet lacp] hierarchy level. To configure LACP system priority for a particular interface, include the system-priority statement at the [edit interfaces ae*x* aggregated-ether-options lacp] hierarchy level. To configure LACP port priority for a particular interface, include the new port-priority statement at the [edit interfaces *interface-name* (fastether-options | gigether-options) 802.3ad ae*x* lacp] hierarchy level.

In addition, the new link-protection statement specifies revertive or nonrevertive protection for aggregated Ethernet interfaces. In revertive mode, when a new link becomes operational or is added to the aggregator, LACP recalculates link priority and the link with higher priority (lower numerical value) becomes the active link. In nonrevertive mode, when collection distribution is enabled, the current active link retains that role even if its priority is lower (is numerically higher) than a subsequently added link. To configure link protection for all aggregated Ethernet interfaces, include the link-protection statement at the [edit chassis aggregated-devices ethernet lacp] hierarchy level. To configure link protection for a particular aggregated Ethernet interface, include the link-protection statement at the [edit interfaces aex aggregated-ether-options lacp] hierarchy level.

The new request lacp link-switchover aex operational mode command supports forced switchover to a different active link. Because the command overrides LACP priority calculations, we strongly recommend that you use it only when the actor (the Juniper Networks router) is controlling the active and standby links and the partner (peer) is following, which applies when you configure only the actor for link protection.

The show lacp interfaces command is enhanced to display link state (active or standby). [*Network Interfaces*, *System Basics*]

- **Extended platform support for Ethernet OAM**—Support for Ethernet Operation, Administration, and Maintenance (OAM) is extended to several types of Gigabit Ethernet PICs installed in Enhanced III FPC3 FPCs in M320 routers and in M120 routers. [*Network Interfaces*]
- **Ethernet host addresses with no subnets**—Enables you to configure an Ethernet interface as a host address (that is, with a network mask of /32), without requiring a subnet. Such interfaces can serve as OSPF point-to-point interfaces, and MPLS is also supported. [*Network Interfaces*, *Routing*]
- **Layer 2 circuit signaling extensions for SAToP and CESoP TDM pseudowires**—Enables wireless service providers to replace T1 and E1 time-division multiplexing (TDM) circuits with T1 and E1 pseudowires for backhauling cell site traffic towards a binary synchronous communications (BSC) device or radio network controller (RNC). Support is added for SAToP or structure-aware TDM circuit emulation service over packet (CESoP) pseudowires on the 12-port T1/E1 Circuit Emulation PIC and the 4-port OC3c/STM1 Circuit Emulation PIC channelized for T1/E1 interfaces using LDP-signaled Layer 2 circuits. To configure SAToP or CESoP encapsulation on a T1 or E1 interface for a PE router, include the satop or cesop statement at the [edit interfaces *interface-name* encapsulation] hierarchy level. [*Network Interfaces*]
- **Command for monitoring Packet Forwarding Engine memory usage (M320 and T-series routing platforms)**—The new show pfe resource usage memory command displays static RAM (SRAM) memory usage statistics for the Packet Forwarding Engine. [*System Basics Command Reference*]
- **Annex B support for SONET multiplex section protection (M120 and M320 routers)**—Enables you to configure support for Annex B standards when you enable multiplex section protection (MSP) switching on SONET/SDH interfaces. Include the sonet-options aps annex-b statement at the [edit interfaces so-*fpc/pic/port*] hierarchy level. Use the show aps extensive command to display detailed information about an Annex B configuration. [*Network Interfaces*]

Services Applications

- **VPN aggregation for VoIP calls**—Supports virtual private network (VPN) routing and forwarding (VRF) instances with the Packet Gateway Control Protocol (PGCP). Users on one VPN can now call users on another VPN, across up to 1024 Layer 3 VPNs. Scalability enhancements include a mesh-like VRF configuration that uses only one logical service interface for each VRF, and the ability to assign a pool of logical service interfaces to a service set, which eliminates the previous requirement for large numbers of separately configured service sets and logical service interfaces.

  To configure a pool of logical services interfaces, include statements at the [edit services service-interface-pools] hierarchy level. To associate a pool with a service set, include the next-hop-service service-interface-pool statement at the [edit services service-set *service-set-name*] hierarchy level.

  To display VRF instances on gates or flows for a particular virtual packet gateway (VPG), include the new destination-routing-instance and source-routing-instance options for the show services pgcp gates *gateway-name* and show services pgcp flows *gateway-name* commands. To view linked VPNs on gates and flows, include the extensive option for the commands. [*Multiplay Solutions*, *Services Interfaces*, *System Basics Command Reference*]

- **Protection against PGCP notification "avalanches" (T640 routing node)**—Prevents a packet gateway from sending an excessive number (an *avalanche*) of media inactivity notifications to the packet gateway controller (PGC) in a short period of time, which can severely degrade PGC performance. For example, when an upstream device in the network fails, all existing terminations in the packet gateway by default report media inactivity immediately, so all notifications arrive at the PGC at about the same time. Avalanche protection enables you to regulate the flow of notifications.

  You can implement avalanche protection in one of two ways:

  - As an H.248 property—Include the ip-flow-stop-detection statement at the [edit services pgcp gateway *gateway-name* h248-properties application-data-inactivity-detection] hierarchy level, specifying either immediate or regulated notification. To configure the transmission frequency for regulated notification, include the notification-regulation default *value* statement at the [edit services pgcp gateway *gateway-name* h248-properties notification-behavior] hierarchy level. To display the settings for regulated notification, issue the show services pgcp gateway *gateway-name* h248-properties notify-behavior command.
  - As a limit on the PIC—Include the rate-limit statement at the [edit services pgcp] hierarchy level. This limits the transmission rate for all message types generated on the PIC, not just inactivity notifications.

  You can log all notifications, along with other *observed events*, by including the audit-observed-events-returns statement at the [edit services pgcp gateway *gateway-name* h248-properties] hierarchy level. To include a timestamp on each event, include the event-timestamp-notification statement at the [edit services pgcp gateway *gateway-name* h248-properties] hierarchy level. [*Multiplay Solutions*, *Services Interfaces*, *System Basics Command Reference*]

- **H.248 overload control (M120, M320, and T640 platforms)**—Enables a packet gateway to notify a PGC when it experiences processing overload that might prevent the timely execution of H.248 transactions. The packet gateway sends an overload notification to the PGC when the number of incoming H.248 transactions that are pending in its queue reaches the configured limit, which is defined as a percentage of the queue size. In response, the PGC lowers the *admitted rate* at which transactions are sent to the packet gateway. If the queue becomes 100 percent full, the packet gateway discards new transactions and sends error code #510 (Insufficient resources). To configure the queue usage percentage at which the packet gateway starts sending overload notifications, include the queue-limit-percentage statement at the [edit services pgcp gateway overload-control] hierarchy level. To display the queue usage percentage, issue the show services pgcp active-configuration command. [*Multiplay Solutions*, *Services Interfaces*, *System Basics Command Reference*]

- **RPM TWAMP support on MultiServices PICs**—Adds support for the RPM Two-Way Active Measurement Protocol (TWAMP) on MultiServices PICs running in Layer 2 mode, as defined in Internet draft draft-ietf-ippm-twamp-09.txt, *A Two-way Active Measurement Protocol (TWAMP)*. You can enable TWAMP on all MultiServices PIC models and on all router platforms that support these PICs. To enable TWAMP, include the following statements in the configuration:

  - The server statement at the [edit services rpm twamp] hierarchy level
  - The client-list and port statements at the [edit services rpm twamp server] hierarchy level, to define at least one client address and a TWAMP listening port
  - The twamp-server statement at the [edit interfaces sp-*fpc/pic/port* unit *logical-unit-number*] hierarchy level, to specify the logical interface that provides the TWAMP service

  To monitor TWAMP functionality, issue the show services rpm twamp server connection and show services rpm twamp server session commands. To clear server connections, issue the clear services rpm twamp server connection command. [*Services Interfaces*, *System Basics Command Reference*]

- **IPv6 flow aggregation templates**—Adds support for IPv6 traffic templates for use with version 9 flow aggregation. Previously, only IPv4 and MPLS templates were supported. Include the family inet6 statement at the [edit forwarding-options sampling input] hierarchy level and the new ipv6-template option at the [edit services flow-monitoring version9 template *template-name*] hierarchy level. To filter the IPv6 traffic on a media interface, include the sampling (input | output) statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet6] hierarchy level. [*Services Interfaces*, *Feature Guide*]

- **Enhanced traffic class support for flow monitoring version 9 (M-series and T-series routing platforms)**—Supports IPv6 with version 9 flow monitoring traffic sampling and templates. In addition, adds support for multiple traffic classes for version 9 flow monitoring traffic sampling and templates. To configure IPv6 on version 9 sampling, include the family inet6 statement at the [edit forwarding-options sampling input] hierarchy level. To configure multiple traffic classes, include more than one of the family options (family inet, family inet6, or family mpls) at the [edit forwarding-options sampling input] hierarchy level. To define IPv6 for the version 9 template record, include the ipv6-template statement

at the [edit services flow-monitoring version9 template *template-name*] hierarchy level. *Services Interfaces*, *Feature Guide*]

Subscriber Access Management

- **New Mobile IP home agent (MX-series routers)**—The JUNOS software now supports Mobile IP, a tunneling-based solution that enables an MX-series router on a user's home subnet to track subscribers who roam beyond traditional network boundaries. The roaming subscribers are known as *mobile nodes* and they are tracked by registration with the Mobile IP *home agent*. The home agent forwards IP packets intended for the mobile node to the *care-of address* (CoA) currently being used by the mobile node.

  Mobile IP is useful in environments where mobility is desired and the traditional land line dial-in model does not provide an adequate solution, and in environments where a wireless technology is used. Both GRE and IP-in-IP static tunnels are supported.

  To configure Mobile IP service, you enable it on one or more interfaces and configure attributes of the virtual network on the router that hosts the home agent. Optionally, you can configure dynamic home assignment to redirect mobile node registration requests to a different home agent. You can also specify how the mobile node is authenticated.

  To enable Mobile IP service on one or more interfaces, include the interface names at the new [edit services mobile-ip home-agent enable-service] hierarchy level.

  To configure the virtual network on the home agent, include the virtual-network statement at the [edit services mobile-ip home-agent] hierarchy level. To specify the home agent's IP address, include the home-agent-address *address* statement at the [edit services mobile-ip home-agent virtual-network] hierarchy level. The IP address must be a loopback address. You can also include the following statements at the [edit services mobile-ip home-agent virtual-network home-agent-address *address*] to configure attributes of the home agent:

  - registration-lifetime *number*—Sets the maximum lifetime that the home agent grants to a mobile node registration. If the mobile node requests a lifetime longer than this value, the home agent assigns this lifetime. When the lifetime expires, the home agent removes the binding entry for the mobile node and tears down the session.
  - revocation-required—Controls whether the home agent can revoke a mobile node registration.
  - timestamp-tolerance *number*—Sets the maximum difference that the home agent accepts between its local time and a mobile node timestamp.

  To configure the method that the home agent uses to authenticate mobile nodes, include the order statement at the [edit services mobile-ip authenticate] hierarchy level. Specify the value aaa for RADIUS authentication (the default), or local for local authentication. For more information about RADIUS authentication of mobile nodes, see the next bullet item in this section.

To configure the home agent to redirect a mobile node's registration request to a different home agent, include the nai statement at the [edit services mobile-ip dynamic-home-assignment home-agent] hierarchy level. Specify the network address identifier (NAI) for the mobile node in one of two formats: *hostname@domain-name*.com or *@domain-name*.com. To specify the address of the home agent for that mobile node, include the home-agent *address* statement at the [edit services mobile-ip dynamic-home-assignment home-agent nai *nai*] hierarchy level.

To trace home agent operations for troubleshooting purposes, include the flag home-agent statement at the [edit services mobile-ip traceoptions] hierarchy level.

The following is a sample Mobile IP configuration:

```
[edit services mobile-ip]
home-agent {
    enable-service {
        ge-0/1/1.0;
        ge-0/2/1.0;
    }
    virtual-network {
        home-agent-address 192.168.3.1 {
            registration-lifetime 100;
            registration-revocation required;
            timestamp-tolerance 200;
        }
    }
}
dynamic-home-assignment {
    home-agent {
        nai tsr23@example.com {
            home-agent 10.1.3.1;
        }
    }
}
```

To monitor and manage a Mobile IP home agent, issue the following new commands:

- clear mobile-ip binding (all | ip-address *address* | nai *identifier*)—Clears the binding between the home agent and either all mobile nodes or the mobile node with the specified IP address or NAI.
- clear mobile-ip host (all | ip-address *address* | nai *identifier*)—Deletes either all mobile nodes and their binding information, or the mobile node with the specified IP address or NAI and its binding information.
- show mobile-ip home-agent overview—Displays summary information about the home agent.
- show mobile-ip home-agent bindings (ip-address *address* | nai *identifier* | summary)—Displays information about bindings between a home agent and either all mobile nodes or the mobile node with the specified IP address or NAI.

- show mobile-ip home-agent virtual-network—Displays information about the virtual network used by the mobile nodes.
- show mobile-ip home-agent traffic—Displays traffic statistics associated with the home agent.

[*Subscriber Access*]

- **RADIUS-based authentication of mobile nodes using Juniper Networks VSAs**—Enables the Mobile IP home agent to use RADIUS to authenticate mobile nodes that request registration (see the previous bullet item in this section). The home agent extracts the NAI or the node address from the registration request. The authentication, authorization, and accounting service (AAA) generates a RADIUS Access-Request message with the NAI or address as the User-Name attribute (the address is used only if the NAI is missing from the mobile node's request). AAA sends the Access-Request message to the RADIUS server, which returns Juniper Networks vendor-specific attributes (VSAs) in Access-Accept messages to provide the security association information required to authenticate the mobile node.

  The User-Password attribute is also required for authentication, but the registration request does not include this attribute. When you define AAA parameters, you must set the User-Password attribute to the value juniper.

  The home address to be used by the mobile node is conveyed to the RADIUS server in the Access-Request message by one of two standard RADIUS attributes. You can statically configure the home address in the Framed-IP-Address attribute. Alternatively, you can configure a local address pool from which to allocate home addresses to mobile nodes. In this case, the pool identification is passed in the Framed-Pool attribute. You configure these attributes when you configure AAA parameters at the [edit access] hierarchy level.

  The Mobile IP home agent uses RADIUS authentication by default. To explicitly specify its use, include the order aaa statement at the [edit services mobile-ip authenticate] hierarchy level.

  You can trace Mobile IP authentication events by including the following statements at the [edit services mobile-ip traceoptions] hierarchy level:

  - flag authentication—Trace includes all authentication-specific events.
  - flag session-db—Trace includes all session-specific events.
  - flag subscriber—Trace includes all events involving subscribers.

  Table 1 on page 14 describes the VSAs configured on the RADIUS server for RADIUS-based authentication during Mobile IP registration. The JUNOS software uses the vendor ID assigned to Juniper Networks (vendor ID 4874) by the Internet Assigned Numbers Authority (IANA).

Table 1: VSAs for Mobile IP Authentication

VSA Number	VSA Name	Definition
[26-84]	Mobile-IP-Algorithm	Authentication algorithm (optional)
[26-85]	Mobile-IP-SPI	Security parameter index (mandatory)

Table 1: VSAs for Mobile IP Authentication *(continued)*

[26-86]	Mobile-IP-Key	Security association MD5 key (mandatory)
[26-87]	Mobile-IP-Replay	Replay timestamp (optional)
[26-89]	Mobile-IP-Lifetime	Maximum registration lifetime (optional)

If you do not configure the mandatory VSAs, mobile nodes cannot be authenticated and are not registered with the home agent. If you do not configure an optional VSA, RADIUS uses a locally configured value or a default value if there is no locally configured value. When the mobile node is authenticated, the Mobile IP home agent caches all subscriber-related information that was returned by the RADIUS server. Consequently, subscriber configuration changes made in the RADIUS server after authentication do not take effect until the subscriber logs out. [*Subscriber Access*]

- **DHCP support for dynamic demux interfaces (MX-series routers)**—The Dynamic Host Configuration Protocol (DHCP) local server and DHCP relay agent now support the creation of dynamic demux interfaces. To configure dynamic demux interfaces for the DHCP local server, include the dynamic-profile *profile-name* statement at the [edit system services dhcp-local-server] or [edit system services dhcp-local-server group *group-name*] hierarchy level. To configure them for the DHCP relay agent, include the dynamic-profile *profile-name* statement at the [edit forwarding-options dhcp-relay] or [edit forwarding-options dhcp-relay group *group-name*] hierarchy level.

  For both the DHCP local server and DHCP relay agent, you can specify the primary dynamic profile that is instantiated when the first subscriber logs in. This conserves interfaces in networks where dynamic demux interfaces are used to represent subscribers. To configure a primary dynamic profile, include the use-primary *primary-profile-name* statement at the appropriate hierarchy level: [edit system services dhcp-local-server dynamic-profile *profile-name*], [edit system services dhcp-local-server group *group-name* dynamic-profile *profile-name*], [edit forwarding-options dhcp-relay dynamic-profile *profile-name*], or [edit forwarding-options dhcp-relay group *group-name* dynamic-profile *profile-name*].

  To limit the number of subscribers who can simultaneously log in for a group of interfaces, include the interface-client-limit statement at the [edit system services dhcp-local-server group *group-name* overrides] hierarchy level for the DHCP local server, and at the [edit forwarding-options dhcp-relay-agent group *group-name* overrides] hierarchy level for the DHCP relay agent. When the limit is reached, new discovery messages are discarded and counted. When the number of logged-in subscribers drops below the limit, new subscribers are allowed to log in. [*Subscriber Access*]

- **Dynamic subscriber interfaces for IP demux interfaces (MX-series routers)**—Enables you to configure an IP demux interface to represent a subscriber interface in a dynamic profile. When a subscriber logs in using a DHCP access method, the demux interface is dynamically created. To configure the settings that DHCP applies to the interface when the subscriber logs in, include the demux0 statement at the [edit dynamic-profiles *profile-name* interfaces] hierarchy level. [*Subscriber Access*]

- **MAC address validation for Ethernet and IP demux interfaces (MX-series routers)**—Enables the router to validate incoming packets by verifying that their source is a trusted IP and Ethernet media access control (MAC) address. The validation provides additional security when subscribers access billable services, because the router can discard packets that do not meet the validation requirements, such as packets with spoofed addresses.

  MAC address validation is supported on IP demux interfaces and aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces (with or without virtual LAN [VLAN] tagging) on MX-series routers only. The interfaces can be statically created or dynamically created using a dynamic profile.

  To configure MAC address validation, include the mac-validate (loose | strict) statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family-name*] hierarchy level for static interfaces and the [edit dynamic-profiles interfaces *interface-name* unit *logical-unit-number* family *family-name*] hierarchy level for dynamic interfaces.

  The output from the show interfaces command is enhanced to indicate whether MAC address validation is enabled. [*Subscriber Access*, *Network Interfaces*, *Interfaces Command Reference*]

- **Optional disabling of automatic ARP table population (MX-series routers)**—Enables you to hide subscriber MAC address information from devices located outside the trusted access domain (for example, an external digital subscriber line access multiplexer [DSLAM]). By default, as communication is established with subscriber-access clients, the router's Address Resolution Protocol (ARP) table is populated with client MAC address and newly-provisioned IP address information that is obtained from the DHCP local server or relay agent. To instead populate the table with less specific information that is safe to share with devices outside a trusted domain, include the no-arp statement at the appropriate hierarchy level:

  - [edit forwarding-options dhcp-relay overrides]
  - [edit forwarding-options dhcp-relay group *group-name* overrides]
  - [edit system services dhcp-local-server overrides]
  - [edit system services dhcp-local-server group *group-name* overrides]

  [*Subscriber Access*]

- **Command to monitor active subscribers**—The new show subscribers command displays information about active subscribers, including subscriber type, username, IP address, dynamic profile name, MAC address, and login time. [*System Basics Command Reference*]

Layer 2 Ethernet Services

- **Port-based access control (MX-series routers)**—Implements the Institute of Electrical and Electronics Engineers (IEEE) standard 802.1x, *Port-Based Network Access Control*, which provides authenticated access to specific router ports. The only unauthenticated traffic accepted on a protected port is 802.1x control packets, which are forwarded to the Routing Engine for processing. Several 802.1x-compliant authentication methods are supported, including RADIUS and Microsoft Active Directory server.

  You can enable port-based authentication on bridged ports, but not on routed ports. Dynamic changes to a user session are supported, which enables you to terminate an authenticated session by using the RADIUS disconnect message defined in RFC 3576, *Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)*.

  To configure port-based authentication, include the authenticator statement at the [edit protocols dot1x] hierarchy level. [*Network Interfaces*]

Routing Policy and Firewall Filters

- **Aggregate policers for different family address types on the same interface**—Enables you to configure separate firewall filters for different family address types that share the same interface and to configure the same policer as an action for each filter. To configure the aggregate policer, include the logical-interface-policer statement at the [edit firewall policer *policer-name*] hierarchy level. To apply the policer as a firewall filter action, include the *policer-name* statement at the [edit firewall family *family-name* filter *filter-name* term *term-name* then] hierarchy level. [*Policy*]
- **Lower minimum policer bandwidth limit (M120, M320, and MX-series routers)**—Enables you to specify a policer bandwidth limit as low as 8000 bits per second (bps) on these routers. On all other routing platforms, the minimum policer bandwidth remains at 32,000 bps. Include the bandwidth-limit *limit* statement at the [edit firewall policer *policer-name* if-exceeding] hierarchy level. [*Policy*]

Routing Protocols

- **BFD protocol support for OSPFv3**—Enables OSPFv3 to use the Bidirectional Forwarding Detection (BFD) protocol to detect network failures for IP version 6 (IPv6) traffic. (OSPFv2 continues to support BFD, as in previous releases.) Failures are detected more quickly than with OSPFv3's default failure detection mechanism, because BFD's failure-detection time limits are shorter. To enable BFD failure detection for OSPFv3, include the bfd-liveness-detection statement at the [edit protocols ospf3 area *area-id* interface *interface-name*] or [edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*] hierarchy level. BFD for OSPFv3 is also supported for logical systems. The show bfd session detail command has been enhanced to display OSPFv3 information. [*Routing*, *Routing Command Reference*]

- **Priority assignment for prefixes in OSPF import policies**—Enables you to assign one of three priority levels (high, medium, and low) to a prefix included in an OSPF import policy. The configured priorities determine the order in which OSPF routes are updated in the routing table after a network topology change, which is useful in a network with a large number of OSPF routes. In general, routes to which you do not explicitly assign a priority are treated as medium priority.

  To configure prefix priority, include the priority (high | medium | low) statement at the [edit policy-options policy-statement *policy-name* term *term-name* then] hierarchy level. To apply the routing policy, include the import statement at the [edit protocols (ospf | ospf3)] or [edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)] hierarchy level. The import policy can also be applied within a routing instance or logical system. The show ospf route detail command has been enhanced to display the priority of network routes. [*Policy*, *Routing*, *Routing Command Reference*]

- **Advertisement of the best external BGP path to internal peers**—Enables a BGP peer to advertise the external path that has the highest local preference value to its internal peers (that is, into an internal BGP mesh group, route reflector cluster, or autonomous system [AS] confederation) even when the overall best path is an internal route. Being able to advertise this external path can be helpful in scenarios known to result in internal BGP route oscillation. To configure, include the advertise-external statement at the [edit protocols bgp group *group-name*] hierarchy level.

  You can also configure BGP to restrict advertisement of the best external path over an active path with an equal cost until the route selection process reaches the step where the multiple exit discriminator (MED) metric is evaluated. As a result, an external path with an AS path worse than that of the active path is not advertised. To configure, include the conditional statement along with the advertise-external statement at the [edit protocols bgp group *group-name*] hierarchy level.

  To configure a BGP export policy to match on routes advertised through this feature, include the state (active | inactive) statement at the [edit policy-options policy-statement *statement-name* term *term-name* from] hierarchy level. [*Policy*, *Routing*]

- **Support for SEND**—Enables you to secure neighbor-discovery messages with the Secure Neighbor Discovery (SEND) protocol, which uses cryptographically generated addresses (CGAs) as defined in RFC 3972, *Cryptographically Generated*

Addresses. SEND is useful in environments where physical security is not assured on the shared link over which IPv6 nodes send neighbor discovery messages to locate one another, determine link-layer addresses, and maintain reachability information about paths to active neighbors.

By default, routers send and receive both secured and unsecured neighbor-discovery messages. To permit acceptance of secured messages only, include the secure-messages-only statement at the [edit protocols neighbor-discovery secure security-level] hierarchy level.

You must also enable the use of CGAs by including the cryptographic-address statement at the [edit protocols neighbor-discovery secure] hierarchy level. Optionally, you can specify a file that contains a public-private key pair and the length of the key by including the key-pair *filename* and key-length *number* statements at the [edit protocols neighbor-discovery secure cryptographic-address] hierarchy level.

To configure timestamp options that help protect against replay attacks, include the timestamp statement at the [edit protocols neighbor-discovery secure] hierarchy level.

You can configure SEND for individual logical systems by including the appropriate statements at the [edit logical-systems *logical-system-name* protocols neighbor-discovery secure] hierarchy level.

The new Secure field in the output of the show ipv6 neighbors command displays SEND information. [*Routing*, *Routing Command Reference*]

Multicast

- **Full support for IGMPv3 and MLDv2**—Enhances the JUNOS implementation of Internet Group Management Protocol version 3 (IGMPv3) and Multicast Listener Discovery protocol version 2 (MLDv2) to support all protocol features. In particular, the JUNOS software now supports exclude mode, which enables a host to specify sources that do not need to send traffic to it. In JUNOS Release 9.2 and earlier, only include mode is supported.

  In the output from the show igmp group and show mld group commands, the new Group mode field reports the mode as Exclude or Include. [*Multicast*, *Routing Command Reference*]

MPLS Applications

- **MPLS applications over non-MPLS enabled networks (M7i, M10i, and M40e routers)**—Enables MPLS to run on networks that do not have MPLS enabled on their core routers, by replacing the top label of the MPLS label stack with an IP-based encapsulation. There are no new statements associated with this feature. To tunnel MPLS traffic through a non-MPLS enabled network, include the tunnel statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. Only IP version 4 (IPv4) tunnels are supported. [*MPLS*]

Routing Policy and Firewall Filters

- **Additional numeric-range match conditions in firewall filters (MX-series routers)**—Extends the set of match conditions for firewall filters on MX-series routers to include the following:

  - (learn-vlan-1p-priority | learn-vlan-1p-priority-except)—Learned IEEE 802.1p VLAN priority
  - (loss-priority | loss-priority-except)—Packet loss priority (PLP)
  - (user-vlan-1p-priority | user-vlan-1p-priority-except)—User IEEE 802.1p VLAN priority

  The four new match conditions based on IEEE 802.1p VLAN priority are valid for the bridging and virtual private LAN service (VPLS) protocol families only (you can include them at the [edit firewall family (bridge | vpls) filter *filter-name* term *term-name* from] hierarchy level).

  The two new match conditions based on PLP are valid for all protocol families only (you can include them at the [edit firewall family (any | bridge | ccc | inet | inet6 | mpls | vpls) filter *filter-name* term *term-name* from] hierarchy level). [*Policy*, *MX-series Solutions*]

- **Port mirroring of Layer 2 bridge and VPLS traffic (MX-series routers)**—Enables you to configure port mirroring for Layer 2 traffic. Include the family (inet | inet6 | vpls) statement at the [edit forwarding-options port-mirroring] hierarchy level. You can configure various types of port mirroring by including the indicated statements:

- To configure multiple instances of port mirroring, include the instance *instance-name* statement at the [edit forwarding-options port-mirroring] hierarchy level.
- To associate a specific port-mirroring instance with a specific DPC, include the port-mirror-instances *instance-name* statement at the [edit chassis fpc *slot-number*] hierarchy level.
- To configure the router to mirror traffic only once, include the mirror-once statement at the [edit forwarding-options port-mirroring] hierarchy level. This option is useful when you are performing port mirroring on both the ingress and egress interfaces.

The extended support of port mirroring means you can now include the port-mirror action at the [edit firewall family (bridge | vpls) filter *filter-name* term *term-name* then] hierarchy level. [*Layer 2*, *System Basics*]

- **Firewall filters within logical systems**—Enables you to configure firewall filters that apply only within a particular logical system by including statements at the [edit logical-systems *logical-system-name* firewall] hierarchy level. Most of the statements that are valid at the global [edit firewall] hierarchy level are also valid for individual logical systems; for details, see the *Junos OS Policy Framework Configuration Guide*.

  In general, the firewall filters for a logical system must be complete and self-contained. They cannot reference firewall elements configured at the [edit firewall] hierarchy level or at the [edit logical-systems *logical-system-name* firewall] hierarchy level for a different logical system.

  If you do not configure firewall filters for a logical system, the filters defined at the [edit firewall] hierarchy level apply to it. If you do configure filters for a logical system, the globally defined filters do not apply. If you want globally defined filters to apply to the logical system, you must explicitly configure them at the appropriate [edit logical-systems *logical-system-name* firewall] hierarchy level. To reduce duplication of effort in this case, define the filters in one or more configuration groups and apply the groups at the appropriate hierarchy levels. [*Policy*]

VPNs

- **Associating VPLS CE-facing interfaces with CE mesh groups**—The JUNOS software now associates VPLS customer edge (CE)-facing interfaces with the CE mesh groups to which they belong, providing better control over the flooding behavior of VPLS routing instances. In JUNOS Release 9.2 and earlier, all CE-facing interfaces are associated with the default CE mesh group, which causes redundant packet forwarding in fully meshed CE devices, because packets arriving at a CE-facing interface are flooded to all other CE-facing interfaces in the default CE mesh group. To associate an interface with a mesh group, include the interface *interface-name* statement in the configuration for the mesh group. [*VPNs*]
- **Clearing MAC addresses for better convergence**—Enables you to remove dynamically learned MAC addresses from the MAC address database, which promotes faster MAC address convergence. To enable MAC address clearing globally, include the mac-tlv-recv and mac-tlv-send statements at the [edit routing-instances *routing-instance-name* protocols vpls] hierarchy level. To enable

MAC address clearing on the routers in a specific mesh group, include the mac-tlv-recv and mac-tlv-send statements at the [edit routing-instances *routing-instance-name* protocols vpls mesh-group *mesh-group-name*] hierarchy level. [*VPNs*]

- **Inter-AS VPLS with MAC processing at the AS boundary router**—Enables a service provider to interconnect customer sites located in different ASs. Inter-AS VPLS requires internal BGP (IBGP) peering between the provider edge (PE) routers within an AS (including the AS boundary routers), and also requires external BGP (EBGP) peering between the AS boundary routers of different ASs. Both a full mesh of multihop EBGP sessions and a full mesh of label-switched paths (LSPs) are required between all AS boundary routers that peer with the AS boundary routers in the other ASs and provide VPLS service (BGP LSPs are sufficient). The route targets configured for the customer's VPLS routing instances across multiple ASs must be identical.

  To configure inter-AS VPLS, include the peer-as all statement at the [edit routing-instances *routing-instance-name* mesh-group *mesh-group-name*] hierarchy level on provider edge (PE) router that participates in the mesh group. Only one mesh group can be configured, and automatic site IDs are not supported. To verify and troubleshoot inter-AS VPLS configurations, issue the show vpls connections command. [*VPNs*, *Feature Guide*]

- **Multiple logical trunk interfaces per physical interface**—Enables you to configure multiple logical trunk interfaces on a single physical interface. To configure, include the vlan-tagging or flexible-vlan-tagging statement at the [edit interfaces *interface-name*] hierarchy level. To configure the list of acceptable VLAN IDs, include the vlan-id-list statement at the [edit interfaces *interface-name* unit *logical-unit-number* family bridge] hierarchy level.

  Including the family bridge statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level enables the logical interface to belong to multiple bridge domains. To designate a logical interface as an automatic member of all bridge domains specified at the [edit interfaces *interface-name* unit *logical-unit-number* family bridge] hierarchy level, include the interface in the virtual switch configuration for the routing instance. To configure a virtual-switch routing instance, include the instance-type virtual-switch statement at the [edit routing-instances *routing-instance-name*] hierarchy level.

  Including the family bridge statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level further enables the logical interface to handle VPLS routing instance traffic but be a part of a virtual switch routing instance. This enables you to configure multiple customer sites to use the same pseudowire, relying on the inner packet's VLAN ID to identify the customer traffic.

  To configure a dual tagged trunk interface, include the flexible-vlan-tagging statement at the [edit interfaces *interface-name*] hierarchy level. Configure the outer tag by including the vlan-tags outer statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level, and configure the inner tag by including the inner-vlan-id-list statement at the [edit interfaces *interface-name* unit *logical-unit-number* family bridge] hierarchy level. [*Layer 2*]

- **VPLS trunk pseudowire interfaces (MX-series routers)**—Enables VPLS functionality within a virtual-switch routing instance (one for which the configuration includes the instance-type virtual-switch statement at the [edit routing-instances *routing-instance-name*] hierarchy level). The interfaces in a

virtual-switch routing instance (those specified by interface *interface-name* statements at the [edit routing-instances *routing-instance-name*] hierarchy level) are Layer 2 trunk interfaces. The VPLS pseudowires in a virtual-switch routing instance (which are created as a result of including the protocols vpls subhierarchy at the [edit routing-instances *routing-instance-name*] hierarchy level for the virtual-switch routing instance) are also trunk interfaces; that is, they belong to all bridge domains in the routing instance and participate in the passing of traffic within each bridge domain. On MX-series routers running JUNOS Release 9.2 and earlier, only VPLS routing instances (those configured with the instance-type vpls statement) support VPLS functionality. [*Layer 2*, *VPNs*]

High Availability

- **Unified ISSU support for additional hardware**—Unified in-service software upgrade (ISSU) supports the following additional platforms, PICs, and FPCs:
  - 1-port OC48c/STM16 ATM2 IQ PICs, which are supported on the M40e, M120, M320, and T-series routing platforms.
  - T1600 Enhanced Scaling FPC4 on the T1600 routing node.
  - TX Matrix platform. Unified ISSU is supported for the same protocols and hardware components as on T-series routing platforms.
  - MX-series routers. You must also enable nonstop active bridging by including the nonstop-bridging statement at the [edit protocols layer2-control] hierarchy level (for more information about prerequisites, see the *Junos OS High Availability Configuration Guide*).

    Unified ISSU does not support the following features and protocols on MX-series routers:
    - Connectivity-fault management (CFM) for Ethernet interfaces, as defined by the IEEE 802.1ag OAM standard
    - Link-fault management (LFM) for Ethernet interfaces, as defined by the IEEE 802.3ah OAM standard
    - LACP

  [*High Availability*]
- **Unified ISSU support for additional protocols**—Unified ISSU now supports the following transport mechanisms and protocols:
  - Layer 2 circuits.
  - Layer 2 Control Protocol (L2CP). To enable ISSU support, you must include the delegate-processing statement at the [edit routing-options ppm] hierarchy level, which enables distributed periodic packet management (PPM) on the Packet Forwarding Engine instead of the Routing Engine. PPM is used to send and receive the hello bridge protocol data units (BPDUs) that maintain L2CP adjacencies. Performing PPM on the Packet Forwarding Engine means that adjacencies are maintained as a new Routing Engine assumes mastership during ISSU, which eliminates disruption in the Layer 2 control plane and minimizes disruption in the Layer 2 data plane.

- LDP-based VPLS.
- Protocol Independent Multicast (PIM), for IPv4 traffic only.

[*High Availability*]

- **NSR support for additional protocols**—Nonstop active routing (NSR) now supports the following additional transport mechanisms and protocols:
  - Layer 2 circuits
  - LDP-based VPLS

  To configure NSR for these protocols, include the same statements in the configuration as for other protocols: the nonstop-routing statement at the [edit routing-options] hierarchy level and the graceful-restart statement at the [edit chassis redundancy] hierarchy level. [*High Availability*, *Multicast*]
- **NSR support for PIM for IPv4**—NSR now supports PIM for IPv4 traffic. State information replicated on the backup Routing Engine includes information about neighbor relationships, join and prune events, rendezvous point (RP) sets, synchronization between routes and next hops, and the forwarding state between the two Routing Engines.

  To configure NSR for PIM, include the same statements in the configuration as for other protocols: the nonstop-routing statement at the [edit routing-options] hierarchy level and the graceful-restart statement at the [edit chassis redundancy] hierarchy level. To trace PIM NSR events, include the flag nsr-synchronization statement at the [edit protocols pim traceoptions] hierarchy level.

  In JUNOS Release 9.3, NSR support varies for different PIM features. The features fall into the three categories described in the following list: supported features, unsupported features, and incompatible features.

  PIM features supported with NSR in JUNOS Release 9.3

  - Auto-RP
  - Bidirectional Forwarding Detection (BFD)
  - Bootstrap router
  - Dense mode
  - Sparse mode (except for some subordinate features mentioned in the following list of unsupported features)
  - Source-specific multicast (SSM)
  - Static RPs

PIM features not supported with NSR in JUNOS Release 9.3

You can configure the following PIM features on a router along with NSR, but during Routing Engine switchover and other outages, their state information is not preserved and traffic loss is to be expected.

- Internet Group Management Protocol (IGMP) exclude mode
- IGMP snooping
- PIM for IPv6 and related features such as embedded RP and MLD
- Policy features such as bootstrap router export and import policies, flow maps, import join policy, neighbor policy, reverse path forwarding (RPF) check policies, RP/DR register message filtering, and scope policy
- Upstream assert synchronization

PIM features incompatible with NSR in JUNOS Release 9.3

NSR does not support the following features in JUNOS Release 9.3, and you cannot configure them on a router enabled for PIM NSR. The commit operation fails if the configuration includes both NSR and one or more of these features:

- Anycast RP
- Draft-Rosen multicast VPNs (MVPNs)
- Local RP
- Next-generation MVPNs with PIM provider tunnels
- PIM join load balancing

JUNOS Release 9.3 introduces a configuration statement that disables NSR for PIM only, so that you can activate incompatible PIM features and continue to use NSR for the other protocols on the router. Before activating an incompatible PIM feature, include the new nonstop-routing disable statement at the [edit protocols pim] hierarchy level. Note that in this case NSR is disabled for all PIM features, not only incompatible features.

NOTE: If both an incompatible PIM feature and NSR are enabled on a router running JUNOS Release 9.2 or earlier, you must perform additional steps when upgrading the router to JUNOS Release 9.3. For more information, see “Upgrading to Release 9.3 on a Router Enabled for Both PIM and NSR” on page 93.

[*High Availability*, *Multicast*]

Class of Service

- **Extended PIC support for per-unit schedulers**—Enables you to configure per-unit schedulers for T1/*N*xDS0 interfaces on the channelized DS3 IQ PIC, and for DS0 interfaces on the channelized STM1 IQ PIC (that is, you can include the per-unit-scheduler statement at the [edit interfaces *interface-name*] hierarchy level). When per-unit schedulers are configured, you can define dedicated schedulers for the logical interfaces on the PIC by including the scheduler-map statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. [*CoS*]

- **CoS functionality on the Enhanced IQ PIC (M40e, M120, M320, and T-series routing platforms)**—The Enhanced IQ PIC provides the following class-of-service (CoS) features:

  - Transmission rate limiting, to prevent low-latency traffic (such as voice) on a queue configured for strict-high priority from starving queues that are serving lower priority packets. Include the rate-limit option to the transmit-rate statement at the [edit class-of-service schedulers *scheduler-name*] hierarchy level. Although intended primarily for queues that serve low-latency traffic, you can rate limit any queue.

  - Substitution of user-defined values for the type-of-service (ToS) bit in incoming packets. To define the values to substitute, include the translation-table statement at the [edit class-of-service] hierarchy level. To apply the values to a logical interface, include the translation-table statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level. You can translate IPv4 ToS values to IPv4 values (DiffServ code point [DSCP] or INET), IPv6 to IPv6 (DSCP), or MPLS to MPLS (EXP bits).

  - Hierarchical scheduling based on configuring a committed information rate (CIR) and peak information rate (PIR) for physical interfaces combined with other parameters for logical interfaces. For example, you can configure the CIR and PIR for a physical interface (internal node) by including the guaranteed-rate and shaping-rate statements respectively, and configure a packet loss priority and guaranteed rate (transmit rate) for the logical interface (leaf node).

  [*CoS*]

- **Ingress BA classification on the Enhanced IQ PIC (M120 routers)**—The Enhanced IQ PIC for the M120 router supports ingress behavior aggregate (BA) classification for DSCP, IP precedence, or EXP bits. The BA classifier sets the forwarding class and PLP bits for the packet and this information is used to forward the packet. [*CoS*]

- **Support for CoS service updates in dynamic profiles (MX-series routers)**—Enables subscribers to merge services when changing services, instead of replacing services as in previous releases. To enable subscribers to change services by merging, include the new variables statement at the [edit dynamic-profiles *profile-name*] hierarchy level. To configure the CoS features for the dynamic profile, include statements at the [edit dynamic-profiles *profile-name* class-of-service] hierarchy level; the valid statements include delay-buffer-rate, guaranteed-rate, scheduler-map, and shaping-rate. The values of the variables are

applied to a subscriber's configuration when authenticating using RADIUS. [*Subscriber Access*]

- **Excess bandwidth and rate limiting on MultiServices PICs (M-series and T-series routing platforms)**—On all versions (100, 300, and 500) of the MultiServices PIC, you can now limit the transmission rate for a link services IQ (lsq-) logical interface in the same way as on other types of queuing PICs. As with the other PICs, the strict-high queue (often assigned to voice traffic) can starve lower priority queues. We recommend that you rate limit the strict-high queue to prevent this, by including the rate-limit option for the transmit-rate statement at the [edit class-of-service schedulers *scheduler-name*] hierarchy level.

  You can also assign a percentage of excess bandwidth to logical interfaces on MultiServices PICs as on other PIC types. Include the excess-rate statement at the [edit class-of-service schedulers *scheduler-name*] hierarchy level.

  Both the transmit-rate and excess-rate statements apply to egress traffic and per-unit schedulers only. Hierarchical and shared schedulers are not supported. To apply these features to an interface, you must also associate the scheduler with a scheduler map and apply the map to the interface configuration at the [edit class-of-service interfaces *interface-name*] hierarchy level. [*CoS*]

- **Ingress DSCP bits for multicast traffic over Layer 3 VPNs (M120, M320, MX-series, and T-series routing platforms)**—Enables you to configure ToS rewrite rules that an ingress provider edge (PE) router applies to the DSCP bits of GRE packets as part of the implementation of the service provider's overall class-of-service policy. Define the rule by including the rewrite-rules statement at the [edit class-of-service] hierarchy level, and apply the rule to the interface by including the rewrite-rules statement at the [edit class-of-service interfaces *interface-name*] hierarchy level. The rewrite rules are applied to all unicast packets and multicast groups.

  Restrictions on this type of rewrite rule include the following:

  - You cannot define different rewrite rules for different multicast groups.
  - IPv6 multicast is not supported, so you cannot define a rewrite rule for DSCPv6 bits.
  - You cannot use rewrite rules of this type to perform EXP marking.

  [*CoS*]

- **BA classification for VPLS based on IEEE 802.1p bits (MX-series, M120, and M320 routers)**—Enables you to configure BA classifiers based on IEEE 802.1p bits. Include the ieee-802.1 *ieee-classifier-name* encapsulated vlan-tag (inner | outer) statement at the [edit class-of-service routing-instances *routing-instance-name* classifiers] hierarchy level. Specify the inner option to base the BA classification on the inner VLAN tag from the inner Ethernet header, or the outer option to base it on the outer VLAN tag from the inner Ethernet header. BA classification based on the outer Ethernet header bits is not supported. [*CoS*]

JUNOS XML API and Scripting

- **Sequential RPCs in automation scripts**—You can now execute sequential remote procedure calls (RPCs) within an automation script. [*Automation*]

- **New JUNOS XML API operational request tag elements**—Table 2 on page 28 lists the JUNOS Extensible Markup Language (XML) operational request tag elements that are new in JUNOS Release 9.3, along with the corresponding CLI command and response tag element for each one.

Table 2: JUNOS XML Tag Elements and CLI Command Equivalents New in JUNOS Release 9.3

Request Tag Element	CLI Command	Response Tag Element
<clear-binding-all>	clear mobile-ip binding all	None
<clear-binding-ip>	clear mobile-ip binding ip-address	None
<clear-binding-nai>	clear mobile-ip binding nai	None
<clear-database-replication -statistics-information>	clear database-replication statistics	None
<clear-dot1x-interface-session>	clear dot1x interface	<dot1x-interface-session>
<clear-dot1x-mac-session>	clear dot1x mac-address	<dot1x-mac-session>
<clear-service-pgcp-statistics>	clear services pgcp statistics	<service-pgcp-statistics-drain-information>
<clear-visitor-all>	clear mobile-ip visitor all	None
<clear-visitor-ip>	clear mobile-ip visitor ip-address	None
<clear-visitor-nai>	clear mobile-ip visitor nai	None
<get-cos-translation-table-information>	show class-of-service forwarding-table translation-table	<cos-translation-table-information>
<get-cos-translation-table-map-information>	show class-of-service translation-table	<cos-translation-table-map-information>
<get-cos-translation-table -mapping-information>	show class-of-service forwarding-table translation-table mapping	<cos-translation-table-mapping-information>
<get-database-replication -statistics-information>	show database-replication statistics	<database-replication-statistics-information>
<get-database-replication -summary-information>	show database-replication summary	<database-replication-summary-information>
<get-dot1x-authentication-failed-users>	show dot1x authentication-failed-users	<dot1x-authentication-failed-users>
<get-dot1x-interface-information>	show dot1x interface	<dot1x-interface-information>
<get-dot1x-interface-mac-addresses>	show dot1x static-mac-address interface	<dot1x-interface-mac-addresses>
<get-dot1x-static-mac-addresess>	show dot1x static-mac-address	<dot1x-static-mac-addresses>
<get-environment-psu-information>	show chassis environment psu	<environment-component-information>
<get-environment-psu-information>	show chassis environment psu	<environment-component-information>

Table 2: JUNOS XML Tag Elements and CLI Command Equivalents New in JUNOS Release 9.3 *(continued)*

Request Tag Element	CLI Command	Response Tag Element
<get-ioc-npc-connectivity-information>	show chassis ioc-npc-connectivity	<ioc-npc-connectivity>
<get-ip-mip-binding-information>	show mobile-ip home-agent binding ip-address	<mip-binding-information>
<get-mip-binding-information>	show mobile-ip home-agent binding	<mip-binding-information>
<get-mip-ha-overview-information>	show mobile-ip home-agent overview	<mip-ha-overview-information>
<get-mip-ha-traffic-information>	show mobile-ip home-agent traffic	<mip-home-agent-traffic-information>
<get-mip-ha-virtual-network-information>	show mobile-ip home-agent virtual-network	<mip-ha-virtual-network-information>
<get-nai-mip-binding-information>	show mobile-ip home-agent binding nai	<mip-binding-information>
<get-subscribers>	show subscribers	<subscriber>
<get-summary-mip-binding-information>	show mobile-ip home-agent binding summary	<mip-binding-information>
<get-system-resource-cleanup-processes-information>	show system resource-cleanup processes	<system-resource-cleanup-processes-information>
<get-vcpu-information>	show chassis vcpu	<vcpu-information>
<set-logical-router>	set cli logical-system	None

[*JUNOS XML API Operational Reference*]

System Logging

- **New and deprecated system log tags**—The following sets of system log messages are new in this release:
  - JCS—Messages generated by the Juniper Control System (JCS) process, which provides the user interface to the control processes for the management modules and switch modules in the JCS 1200 platform.
  - LLDPD—Messages generated by the Link Layer Discovery Protocol (LLDP) process, which is responsible for discovering devices on a network and their capabilities.
  - SAVAL—Messages generated by the MAC SA validation process, which verifies the match between the source IP address and MAC address in incoming packets, to prevent spoofing on Ethernet-based interfaces.
  - VCCPD—Messages generated by the Virtual Chassis Control Protocol (VCCP) process.

The following system log messages are new in this release:

- CHASSISD_PSU_ERROR
- CHASSISD_PSU_FAN_FAIL
- CHASSISD_PSU_INPUT_BAD
- CHASSISD_PSU_OVERLOAD
- CHASSISD_PSU_TEMPERATURE
- CHASSISD_PSU_VOLTAGE
- CHASSISD_SNMP_TRAP1
- DCD_PARSE_ERROR_SCHEDULER_LIMIT
- L2ALD_MAC_LIMIT_REACHED_IFBD
- LIBMSPRPC_CLIENT_KCOM_NO_IF
- SPD_GEN_NUM_FAIL

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:

- KMD_CFG_NO_TRACE_FILE
- KMD_VPN_BIND_TUNNEL_IF
- PFE_FW_UNSUPPORTED

[*System Log*]

Related Topics

- Changes in Default Behavior and Syntax in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 31
- Issues in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 38
- Errata and Changes in Documentation for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 87
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 90

Changes in Default Behavior and Syntax in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms

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.

Platform and Infrastructure

- **Configuring TACACS+ accounting**—Two new statements have been introduced at the [edit system tacplus-options] hierarchy level to support the logging of accounting records to the correct log file on a TACACS+ server. When you include the no-cmd-attribute-value statement, the cmd attribute value is set to an empty string in the start and stop requests for accounting of login events. When you include the exclude-cmd-attribute statement, the cmd attribute is omitted from the start and stop requests for accounting of login events. [*System Basics*]
- **New range for ARP aging timer**—The range for the aging-timer statement at the [edit system arp] hierarchy level is now 5 through 240 minutes, instead of 20 through 240 minutes as in JUNOS Release 9.2 and earlier. [*System Basics*]

User Interface and Configuration

- **Deprecated trace options**—For M-series, MX-series, and T-series routing platforms, the no-stamp and replace statements have been deprecated at the [edit *any-level* traceoptions file] hierarchy level (in other words, at all hierarchy levels at which the traceoptions statement is supported). [All JUNOS documentation for M-series, MX-series, and T-series routing platforms]

Interfaces and Chassis

- **New output field explaining reason for Routing Engine reboot**—A new field Last reboot reason has been added to the output from the show chassis routing-engine operational mode command. This field displays the reason the Routing Engine last rebooted. [*System Basics Command Reference*]
- **Enhanced output from the 'show interfaces' command**—The show interfaces (detail | extensive) commands now display multiple input or output filters, including any filters attached to the interface through dynamic service activation. [*Interfaces Command Reference*]

Services Applications

- **Configuring the preshared key for an IKE policy**—A local certificate is an alternative to the preshared key for an Internet Key Exchange (IKE) policy. The commit operation fails if neither a preshared key nor a local certificate is configured at the [edit security ike policy *peer-address*] hierarchy level. [*System Basics*]
- **Support for H.248 inactivity timer package**—The packet gateway supports the inactivity timer package defined in International Telecommunication Union Telecommunication Standardization (ITU–T) Recommendation H.248.14, *Gateway control protocol: Inactivity timer package* (March 2002). The packet gateway can now detect the failure of its active PGC through message inactivity. The inactivity timer is applied to the root terminations of a VPG. To configure, include statements at the [edit services pgcp gateway *gateway-name* h248-properties inactivity-timer inactivity-timeout] hierarchy level. [*Services Interfaces*]
- **Latch deadlock and media inactivity detection**—Enables you to generate data inactivity notifications for PGCP gates without latching events, extending the previous functionality which applied only to gates with latching events. The feature enables you to specify a different delay (before inactivity measurement begins) for gates with latching events and gates without latching events, and to specify an inactivity timer that applies to either type of gate. You can also specify that a gate experiencing an inactivity delay be forced out of service, in which case the gate discards all packets and stops sending inactivity notifications.

  To configure latch deadlock and media inactivity detection, include the following statements at the [edit services pgcp gateway *gateway-name* data-inactivity-detection] hierarchy level: inactivity-delay, inactivity-duration, latch-deadlock-delay, report-service-change, and stop-detection-on-drop. [*Services Interfaces*]
- **Deprecated PGCP inactivity statements**—As a consequence of the new media inactivity detection feature described in the preceding bullet item, the following

statements at the [edit services pgcp gateway *gateway-name*] hierarchy level are deprecated in JUNOS Release 9.3 and later: gate-inactivity-delay, gate-inactivity-duration, and latch-deadlock-duration. [*Services Interfaces*]

- **Specifying information included in packet gateway messages**—You can now specify which "method" and "reason" values a VPG includes in ServiceChange commands that it sends to the PGC when the state of a control association, virtual interface, or context changes. To configure, include the statements at the [edit services pgcp gateway *gateway-name* h248-options service-change] hierarchy level. [*Services Interfaces*]
- **Changes to PGCP virtual interface configuration**—If you do not require ingress filtering, you no longer need to configure a physical (media) interface in the PGCP virtual interface configuration. In the following example, ingress filtering is applied:

  ```
  virtual-interface 0 {
     media-service mmm;
     interface fe-0/3/3.0;
  }
  ```

  In the following example, ingress filtering is not applied:

  ```
  virtual-interface 0 {
     media-service mmm;
  }
  ```

  Also, the no-ingress-interface-filtering statement is deprecated at the [edit services pgcp virtual-interface *interface-name*] hierarchy level. [*Services Interfaces*]
- **Changes to PGCP statistics reporting**—In JUNOS Release 9.3R2 and later, the output of the show services pgcp statistics gateway command is changed to include columns that report the number of commands received and sent, the number of commands that use a wildcard termination ID, the number of commands generating a "success" reply, and the number of commands generating an "error" reply. [*System Basics Command Reference*]
- **Random port allocation for NAT**—By default, the JUNOS software allocates NAT ports sequentially. To configure random port allocation, include the random-allocation statement at the [edit services nat pool *pool-name* port (automatic | range low *minimum-value* high *maximum-value*)] hierarchy level. [*Services Interfaces*]

Subscriber Access Management

- **Enforcement of license requirements at runtime**—The configuration for some JUNOS features is shared by multiple services. If one service requires a license for the feature and the other does not, the license requirement is enforced when the feature is used, not when the configuration is committed as was previously the case. For example, the AAA and Layer 2 Tunneling Protocol (L2TP) services share the configuration for authentication order, but only AAA requires a license. A warning no longer appears at commit time if one or both AAA and L2TP are configured but the license is not configured; instead, the license requirement is enforced when AAA authenticates a client. [*Subscriber Access*]

Layer 2 Ethernet Services

- **Global configuration of Layer 2 learning properties (MX-series routers)**—To configure Layer 2 learning properties that apply system-wide on an MX-series router, include the global-mac-statistics, global-mac-table-aging-time *seconds*, and global-no-mac-learning statements at the [edit protocols l2-learning] hierarchy level. The mac-statistics, mac-table-aging-time, and no-mac-learning statements have been deprecated at this hierarchy level (but still apply at other levels). The global-mac-limit *number* statement continues to be supported at the [edit protocols l2-learning] hierarchy level. The mac-statistics and no-mac-learning statements continue to be supported at the [edit bridge-domains *domain-name* bridge-options] hierarchy level, and (in JUNOS Release 9.2 and later) at the [edit switch-options] hierarchy level. [*Layer 2*]
- **Maximum number of active logical interfaces (MX-series routers)**—On MX-series routers, a maximum of 4000 active logical interfaces are now supported on a bridge domain or on each mesh group in a VPLS instance configured for Layer 2 bridging. [*Routing*]

Routing Protocols

- **Change to range for link-state PDU interval for IS-IS**—For the lsp-interval *milliseconds* statement at the [edit protocols isis *interface-name*] hierarchy level, the range of valid values is now 0 through 1000. The default value remains 100 milliseconds. [*Routing*
- **Support for IS-IS traffic engineering shortcuts extended to IPv6 routes**—Enables you to configure IS-IS to use interior gateway protocol (IGP) shortcuts for IPv6 routes. Previously, only IPv4 routes were supported. As a result, the hierarchy for configuring IS-IS IGP shortcuts is changed. You can include the new shortcuts statements at the indicated hierarchy levels:

```
[edit protocols isis traffic-engineering]
family inet {
   shortcuts {
      multicast-rpf-routes;
   }
}
family inet6 {
   shortcuts;
```

```
}
```

LSPs to be used for shortcuts continue to be signaled using IPv4. However, by default, shortcut routes calculated through IPv6 routes are added to the inet6.3 routing table. The default behavior is that only BGP uses LSPs in its calculations. If you configure MPLS so that both BGP and IGP use LSPs for forwarding traffic, shortcut routes calculated through IPv6 are added to the inet6.0 routing table.

You can use the legacy configuration at the [edit protocols isis traffic-engineering shortcuts] hierarchy level to enable IPv4 shortcuts and automatically disable IPv6 shortcuts.

In addition, the show isis overview command has been enhanced to display shortcuts for both IPv4 and IPv6. [*Routing*, *Routing Command Reference*]

- **Support for AS-dot notation for 4-byte AS numbers**—Enables you to configure and display an AS number using the AS-dot notation of two integer values joined by a period character: *high-order-16-bit-value-in-decimal*.*low-order-16-bit-value-in-decimal*. The JUNOS software continues to support the plain-number format for 2-byte and 4-byte AS numbers.

  To configure the AS number for the router, include the autonomous-system *as-number* statement at the [edit routing-options] hierarchy level. For *as-number*, specify a value from 0.0 through 65535.65535 (AS-dot notation).

  You can also use the AS-dot notation for other statements that support 4-byte AS numbers. For example, you can also configure a 4-byte AS number using this format for route-target and route-origin BGP extended communities and the route-distinguisher identifier, as well as the BGP peer AS and the BGP local AS.

  By default, operational mode commands display 4-byte AS numbers as plain numbers. To display AS numbers in AS-dot notation, include the asdot-notation statement at the [edit routing-options autonomous-system] hierarchy level. [*Routing*, *Policy*, *Routing Command Reference*]

- **Extended BGP support for vendor-specific outbound route filtering capability codes**—Extends support to BGP groups and neighbors. The bgp-orf-cisco-mode statement enables interoperability with routers that use the outbound route filtering capability code of 130 and the code-type of 128 for prefix-based outbound route filters, instead of the standard code of 3 and code-type of 64. When included at the [edit routing-options] hierarchy level (previously the only valid location), the statement applies to all BGP peers configured on the router. You can now enable interoperability for a particular BGP group or peer by including the bgp-orf-cisco-mode statement at the following hierarchy levels:

  - [edit protocols bgp outbound-route-filter]
  - [edit protocols bgp group *group-name* outbound-route-filter]
  - [edit protocols bgp group *group-name* neighbor *address* outbound-route-filter]

  The show bgp neighbor command is enhanced to display whether interoperability is enabled with routers that use the nonstandard capability codes. [*Routing*, *Routing Command Reference*]

Multicast

- **Some PIM clear commands not supported on backup Routing Engine**—The clear pim join, clear pim register, and clear pim statistics operational mode commands are not supported on the backup Routing Engine of routers running JUNOS Release 9.3 and later. [*Multicast*, *Routing Command Reference*, *Policy*]
- **Changes to IGMP behavior based on version number**—In JUNOS Release 9.1 and later, the IGMP version configured for a particular interface (by including the version statement at the [edit protocols igmp interface *interface-name* hierarchy level) overrides the version configured for all interfaces on a routing platform (by including the version statement at the [edit protocols igmp interface all hierarchy level).

  In JUNOS Release 9.3 and later, if you specify a source address for a static multicast group (by including the source *address* statement at the [edit protocols igmp interface *interface-name* group *group-name*] hierarchy level), you must also set the IGMP version to version 3 by including the version 3 statement at the [edit protocols igmp interface (*interface-name* | all)] hierarchy level. If the IGMP version is not IGMPv3, the specified source is ignored and only the group added. The join is treated as an IGMPv2 group join. [*Multicast*]

VPNs

- **MTU mismatch between PE routers is allowed**—To configure the JUNOS software to allow a Layer 2 circuit to be established even though the maximum transmission unit (MTU) configured on the PE router does not match the MTU configured on the remote PE router, include the ignore-mtu-mismatch statement at the [edit protocols l2circuit neighbor *address* interface *interface-name*] hierarchy level. [*VPNs*]

High Availability

- **VRRP enhancements for ARP requests**—When a router responds to an ARP request, the Virtual Router Redundancy Protocol (VRRP) virtual MAC address is sent as the Ethernet source address in the Ethernet frame. When VRRP and proxy ARP are both configured, only the VRRP master on that subnet responds to proxy ARP requests. In previous software releases, when a router responded to an ARP request, the hardware MAC address was sent as the Ethernet source address. When VRRP and proxy ARP were both configured, the router responded as proxy for an ARP request if the address was reachable, irrespective of the VRRP state (backup or master) on the subnet. [*High Availability*]
- **Corrected interface status information from 'show vrrp' command**—When the VRRP virtual interface is down, the Master router: field of the output from the show vrrp (detail | extensive) commands now reports the value N/A, instead of the IP address of the last known master router as was reported previously. [*Interfaces Command Reference*]
- **Profile information in 'show vrrp' command output**—The delay threshold and computed-send-rate fields in the output of the show vrrp profile statistics command

are now also included in the output of the show vrrp (detail | extensive) commands. [*Interfaces Command Reference*]

- **Shorter minimum interval between successive Routing Engine switchover events**—The minimum interval between successive Routing Engine switchover events is 4 minutes (240 seconds) instead of the previous requirement of 5 minutes (300 seconds). [*High Availability*]

Class of Service

- **Not all filter match conditions are supported by the Enhanced IQ DPC**—On MX-series routers with the EQ DPC, forwarding class is not supported as a match condition in a firewall filter. [*CoS*]

Forwarding and Sampling

- **Firewall filter based on forwarding class for logical systems**—A firewall filter configured for a logical system can now match or set a packet's forwarding class. To configure, include the forwarding-class *class* statement at the following hierarchy levels:
  - [edit logical-systems *logical-system-name* firewall family *family-name* filter *filter-name* term *term-name* from] (to match on forwarding class)
  - [edit logical-systems *logical-system-name* firewall family *family-name* filter *filter-name* term *term-name* then] (to set the forwarding class)

  The specified *class* must be configured at the [edit class-of-service] hierarchy level. As a general rule, firewall configurations defined under logical systems must be self-contained and cannot reference configurations outside the logical system hierarchy. However, this statement is allowed. It facilitates global, router-wide configurations for forwarding classes. [*Policy*]
- **Configurable hash seed for load balancing**—On routing platforms with the Internet Processor II application-specific integrated circuit (ASIC), all Packet Forwarding Engine slots are assigned the same hash seed by default. You can now configure a unique load-balance hash seed for each slot, enabling better utilization of available links. To configure, include the per-flow hash-seed *number* statement at the [edit forwarding-options load-balance] hierarchy level. [*Policy*]

Related Topics

- Features in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 5
- Issues in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 38
- Errata and Changes in Documentation for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 87
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 90

Issues in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms

The current software release is Release 9.3R4. For information about obtaining the software packages, see "Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms" on page 90.

Current Software Release

The current software release is Release 9.3R4. For information about obtaining the software packages, see "Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms" on page 90.

Outstanding Issues

Platform and Infrastructure

- If the tunnel destination is in a VPN, the generic routing encapsulation (GRE) traffic may get deleted due to a lookup in the wrong forwarding table. [PR/45035]
- 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]
- 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]
- If you configure 11 or more logical interfaces in a single VPLS instance, VPLS statistics might not be reported correctly. [PR/65496]
- 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]
- In the situation where a Link Services (LS) interface to a CE router appears in the VPN routing and forwarding table (VRF table) and a 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]
- On the T-series routing platform, when you include the no-labels configuration 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]

- Traceroute does not work when ICMP tunneling is configured. [PR/94310]
- The initialization fails to parse the configuration present in the init.conf file. [PR/94576]
- For T-series and M320 routers, multicast traffic with the do not fragment bit set to a low MTU value is being dropped. If the clear pim join command is executed, the router stops forwarding all traffic transiting the interface. [PR/95272]
- A firewall filter that matches the forwarding class of incoming packets (that is, includes the forwarding-class 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]
- The JUNOS software does not 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, configure the next-hop address as a static ARP entry by including the arp *ip-address* statement at the [edit interfaces *interface-name*] hierarchy level. [PR/237107]
- Currently, the JUNOS Software cannot build an outbound serial connections through the AUX port. For example, build an outbound serial connection to a console on an adjacent router. [PR/256818]
- On T640, T320, and M320 routers, if you take an FPC offline during an ISSU boot, other FPCs in the router might crash. This happens when there is transit traffic flowing from the other FPCs towards the offlined FPC. [PR/268294]
- When Periodic Packet Management (PPM) delegation for Bidirectional Forwarding Detection (BFD) sessions is disabled (the delegate-processing statement is removed at the [edit routing-options ppm] hierarchy level), the BFD sessions might be terminated (because a "state is down" message is sent) and reestablished. [PR/280233]
- When you perform an in-service software upgrade (ISSU) on a routing platform with an FPC3 or an Enhanced FPC3 with 256 MB of memory and the number of routes in the routing table exceeds 750,000, route loss might occur. If route loss occurs, as a workaround, perform either of the following tasks: (a) replace the FPC3 or Enhanced FPC3 with another FPC that has more memory, or (b) after the ISSU is complete, reboot only the FPC3 or Enhanced FPC3. [PR/282146]
- For Routing Engines rated at 850 MHz (which appear as RE-850 in the output from the show chassis hardware command), messages like the following might be written to the system log when you insert a PC Card: “bad Vcc request” and “Device does not support APM.” Despite the messages, operations that involve the PC Card work properly. [PR/293301]
- Next-hop marking (marked with a dash) in the show route forwarding-table command output indicates which next hops might not transmit traffic in a hierarchical load-balancing topology (for example, multiple load-balanced LSPs over multiple paths or aggregated interfaces). The forwarding-options indexed-next-hop statement was added to address hierarchical load-balancing issues, but configuring this statement may result in the next-hop marking being inaccurate and so the markings should be ignored. [PR/293306]

- Temporary files named in the format cprod*xxxxxx* are retained in the temporary directory on the router and can be deleted. [PR/304750]
- On a Protected System Domain, under the following conditions an FPC might generate a core file and stop operating: (a) a firewall policer with a large number of counters (for example, 20,000) is applied to a shared uplink interface and (b) the FPC that houses the interface does not have a sufficiently powerful CPU. As a workaround, reduce the number of counters or install a more powerful FPC. [PR/311906]
- The SSB servers display an error when you delete a string from the redix tree and then reboot. [PR/312453]
- When you commit a configuration that includes the dynamic demux relay feature and there is a large number of subscribers (for example, 64,000), all subscribers do not become active and the kernel generates an error. [PR/312563]
- Traffic originating from a remote PE router is silently dropped without informing the source that the data did not reach its intended recipient when the multicast MAC address is configured on the local PE router for a CE device. [PR/398698]
- Following an FPC reset, the next-hop route pointing to the service PIC interface running RPM might be incorrect. [PR/438599]

User Interface and Configuration

- The CLI does not generate a warning if multiple users are configured with the same user ID. [PR/55774]
- 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]
- The logical system 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 you are working in private configuration mode and try to commit a configuration that includes a comment about an inactive configuration statement, the commit operation fails with the message "syntax error.". [PR/270160]
- In the output from the configuration mode show | compare command, the banner might be the parent level of the current hierarchy level instead of the current level itself. For example, when the current hierarchy level is [edit interfaces fe-1/1/1], the banner in the output reads [edit interfaces], but the additions and deletions are reported with respect to the [edit interfaces fe-1/1/1] level. [PR/291574]
- A user belonging to a login class with limited rights to modify a specific firewall filter cannot use the insert command to reorder firewall terms. [PR/310872]
- The IPv6 PMTU discovery timeout variable is ip6_pmtu_timeout instead of path_mtu_timeout. [PR/315133]

- When executing the commit sync command, messages appear on the backup Routing Engine. These messages can be ignored. [PR/395716]
- Using the filter config text in the NETCONF get-config command results in a syntax error and the router configuration cannot be returned in ASCII format. [PR/430799]

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 channelized E1 interfaces, you might be able to configure clocking on ds-*fpc*/*pic* /*port*:*n* interfaces, where *n* is not unit 0. This is an invalid configuration and might cause a clocking selection problem on the other channels. [PR/24722]
- On a 2-port OC12 ATM2 IQ interface, the total virtual path (VP) downtime might not display correctly in the show interfaces command output. [PR/27128]
- 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* | logical-systems *logical-system-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 the Channelized STM1 with a QPP PIC, error monitoring for CRC and frame errors might not work as expected. [PR/39440]
- 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, 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]

  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 has changed interface MTU settings prior to the switchover event. [PR/61121]
- If you configure graceful Routing Engine switchover (GRES) 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]

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]

- 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]
- XFP-OC192-SR may report "XFP read fail, retry for 1 times" randomly. This is a cosmetic issue and doesn't affect to the interface functionality. [PR/262883]
- The hot swapping fan tray for the M120 might cause the Check CB alarm to activate. [PR/268735]
- On the JCS 1200, when you issue the clear -config -T switch[1] command using the management module, the switch module returns to its factory default setting instead of the Juniper Networks default setting. As a workaround, do not issue the command. [PR/274399]
- When you configure ILMI on an ATM interface (include the ilmi statement at the [edit interfaces *interface-name* atm-options] hierarchy level) and a graceful Routing Engine switchover (GRES) or unified in-service software upgrade (ISSU) event occurs, the show ilmi command no longer returns any output. [PR/282051]
- On a router with Frame Relay multilink configured on a MultiServices 400 PIC or on a channelized DS3 PIC, when the minimum links value for the Frame Relay interface is set to 8 and a link is deactivated from the configuration, the link remains up. [PR/285244]
- On the Juniper Control System (JCS) platform, the control and management traffic for all Routing Engines share the same physical link on the same switch module. In rare cases, the physical link might become oversubscribed, causing the management connection to Protected System Domains (PSDs) to be dropped. [PR/293126]
- On a Protected System Domain (PSD) configured with a large number of BGP peers and routes (for example, 5000 peers and a million routes), FPCs might restart during a graceful Routing Engine switchover. [PR/295464]
- When two routers are connected via SONET/SDH interfaces that are configured as container interfaces and the Routing Engine on one router reboots, the container interfaces on the other router might go down and come up again. [PR/302757]
- On M5, M10, M20, and M40 routers, when you issue an SNMP query for alarm LED status (such as the show snmp mib walk jnxLEDState command), the message “FPM device not open” might be logged. This is an erroneous message and can be ignored. [PR/313073]
- On MX-series routers, the path MTU discovery for a GRE tunnel is not functioning properly. [PR/390993]
- In JUNOS Release 9.3 and later, VPLS customer edge (CE)-facing interfaces can be associated with the CE mesh groups to which they belong, instead of only with the default CE mesh group (as in JUNOS Release 9.2 and earlier). However, the JUNOS Release 9.2 behavior still applies to interfaces in a VPLS routing

instance that is defined at the [edit logical-systems *logical-system-name* routing-instances] hierarchy level. Also, if you move the configuration for a logical interface in a VPLS routing instance from the [edit routing-instances *routing-instance-name*] hierarchy level to the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*] hierarchy level, the value vpls might stop appearing in the Proto column of the output from the show interfaces terse command. As a workaround, perform the move in two steps by removing the interface from the [edit routing-instances *routing-instance-name*] hierarchy level and committing the configuration, then creating the interface at the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*] hierarchy level and committing again. [PR/400248]

- The XML output is not correct when the Virtual Router Redundancy Protocol (VRRP) track interface is configured. [PR/414734]
- On MX-series routers, MAC address accounting in the egress direction might not work if traffic is unidirectional and no traffic flows in the reverse direction for a duration longer than the aging interval. [PR/415146]
- When you configure the payload port-data statement at the [edit family mpls hash-key] hierarchy level on M120, MX-series, or M320 platforms with E3 FPCs, the hashing algorithm might not take the port-data values into account. [PR/442223]

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]
- 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]
- 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]
- 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]
- If a large number of BGP authentication sessions (for example, 400) are configured in a VRF instance, the following message is written to the system log when the configuration is committed: “keyadmin[*pid*]: dump_assn: posting additional read." This message can be ignored and there is no operational impact. [PR/295407]
- A user belonging to a login class with limited rights to modify a specific firewall filter cannot use the insert command to reorder firewall terms. [PR/312961]

- The IPv6 PMTU discovery timeout variable is ip6_pmtu_timeout instead of path_mtu_timeout. [PR/401247]
- As a fix, the Multilink Point-to-Point Protocol (MLPPP) reassembly logic does not perform a strict out-of-order check. In a multi-CPU packet handling environment, packets arriving later may be processed before the first. [PR/430296]

Subscriber Access Management

- When dynamic IP address assignment is configured, if there is only one address left in the address allocation pool and an attempt to authenticate with a service fails (because, for example the authentication request specifies an invalid service name), a subsequent authentication attempt for the service also fails. The following messages might appear in the log for the authentication process (authd): "assigned address *address* in use, trying next available" and "Unable to assign an address." [PR/305516]
- When you use a RADIUS Change-of-Authorization (CoA) message to activate a service that is already activated, the service is removed. [PR/307983]

Routing Policy and Firewall Filters

- On M-series and T-series routers running JUNOS Release 9.3R1 and later, FPCs might stop functioning if you configure a firewall filter and include the family any statement at the [edit firewall] hierarchy level, and apply the filter to an interface for which the configuration includes the family iso statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. Apply a firewall filter that is configured with the family any statement only to an interface that is not configured with the family iso statement. [PR/408617]
- On an MX-series router, if you configure a logical interface policer containing the bandwidth-limit and burst-size-limit statements at the [edit firewall police *policer-name* if-exceeding] hierarchy level, then perform an in-service software upgrade (ISSU) from JUNOS Release 9.3 to 9.4, load another configuration, issue the ping command to verify connectivity to an adjacent neighbor, and finally perform a rollback to the first configuration, you might not be able to reach the neighbor again when you reissue the ping command because the ICMP packet flow might be blocked. [PR/408893]

Routing Protocols

- 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 a router receives a Pragmatic General Multicast (PGM) Source Path Message (SPM), it does not create a forwarding cache, nor does it forward the message to other routers as a heartbeat, as specified in RFC 3208. Also, the routers multicast cache might time out if it does not receive actual PGM data (ODATA) for more than 6 minutes. As a workaround, configure the PGM source application to send PGM ODATA at least once every 6 minutes. The ODATA acts as the heartbeat message in lieu of the SPM messages and ensures that the multicast and forwarding caches are created and updated. [PR/37504]

- 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]
- If a BGP group is created but without any defined peers, a warning message appears when the configuration is committed. [PR/63279]
- 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]
- In the output of the show pim join extensive command, the assert winner status is displayed in the Outgoing Interface List (OIL) for PIM Dense Mode (PIM-DM) but not for auto-RP dense groups. [PR/74737]
- If ICMP tunneling is enabled on the router and you configure a new logical system that does not have ICMP tunneling enabled, the feature is globally disabled. [PR/81884]
- 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]
- When you commit a new configuration for nonstop routing (NSR) on a primary Routing Engine that differs from the configuration for NSR that is already running on the backup Routing Engine, the routing protocol process stops functioning on the backup Routing Engine only. Traffic forwarding is not affected. [PR/254379]
- Disabling the PIM protocol with the set protocols pim disable command can cause the router to stop operating until that statement is removed. As a workaround, use the deactivate protocols pim command instead. [PR/274478]
- The routing protocol process may restart if PIM is configured to run on unnumbered interfaces. [PR/295319]
- The clear ospf io-statistics command may not clear the counter values that would be seen using the show ospf io-statistics command. [PR/308679]
- The clear ospf io-statistics command might not clear the counter values that are displayed by the show ospf io-statistics command. [PR/401351]
- The show isis statistics command does not display the IS-IS packet statistics. [PR/405022]
- OSPF and IS-IS differ in how they handle the addition of a better internal or external (smaller IGP metric) route into the protocol's internal routing-table. IS-IS flushes all next-hops information (including LSP next-hops) when learning a better prefix, despite equal-cost LSP tunnels, whereas OSPF does not. However, this does not cause any issues with respect to load balancing. [PR/408702]
- The rendezvous point (RP) is not learned on a router where auto-rp discovery is configured. A mismatch occurs between the PIM interface configuration on a router where auto-rp discovery is configured and on a router where auto-rp

mapping is configured. For example, one router has an IFL with PIM configured and the other has an IFL with PIM disabled. As a workaround, ensure that PIM is enabled on all IFLs on both routers. [PR/445917]

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 cross-connected circuit (CCC) traverses a forwarding adjacency 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 modify the primary path for an MPLS LSP by using the delete protocols mpls label-switched-path *lsp-path-name* primary *path-name* command in configuration mode, followed by the set protocols mpls label-switched-path *lsp-path-name* primary *path-name* command, and then issue the commit command, the entire LSP (both primary and secondary) is torn down and then rebuilt from scratch. As a workaround, issue the delete protocols mpls label-switched-path *lsp-path-name* primary *path-name* command in configuration mode, followed by the commit command. Then issue the set protocols mpls label-switched-path *lsp-path-name* primary *path-name* command, followed by the commit' command. [PR/62365]
- 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]
- An error in the Constrained Shortest Path First (CSPF) software might cause the routing protocol process (rpd) to generate a core file and stop operating. [PR/103777]
- When there are more than five link-protected or node-link-protected LSPs to the same destination and per-packet load balancing is enabled, some bypass next-hops might not be part of the active route. This can occur after a primary link goes down and comes back up. [PR/259219]
- For point-to-multipoint LSPs configured for VPLS, the ping mpls command reports 100 percent packet loss even though the VPLS connection is active. [PR/287990]
- The monitor label-switched-path output control key "n" does not work. [PR/298814]

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]
- Traffic might not flow when an ATM interface is used as the access circuit on an M120 router. [PR/255160]

- If a PE router is acting as the mapping agent for PIM auto-RP, PR elections based on the bootstrap router (BSR) do not work correctly. [PR/305325]
- For a VRF instance configured for PIM, MVPN, and provider tunnels (the pim and mvpn statements are included at the [edit routing-instances *vpn-name* protocols] hierarchy level and the provider-tunnel statement is included at the [edit routing-instances *vpn-name*] hierarchy level), when PIM is deactivated and reactivated, it fails to install type-5 (source-active) routes in the *instance-name*.mvpn.0 routing table. This issue arises only when remote C-multicast joins are configured on the ingress PE router (as displayed by the show mvpn c-multicast command). [PR/306983]
- When an LSP switches from a primary path to a bypass path, Layer 2 circuits might go down and come up again, resulting in packet loss. [PR/309085]
- In JUNOS Release 9.3, when you configure inter-AS VPLS with MAC processing at the autonomous system (AS) boundary router along with multihoming, and if a designated forwarding AS boundary router fails and then comes back up again, traffic flowing to the local AS from the other AS's boundary router might be lost. The loss occurs in the time period (tenths of a second) during which the old designated forwarding AS boundary router is taking back the role of designated forwarder. [PR/312730]
- On a router configured for NSR (the nonstop-routing statement is included at the [edit routing-options] hierarchy level), if an NSR switchover occurs after the configuration for routing instances changes in certain ways, BGP sessions between PE and CE routers might not be established after the switchover. [PR/399275]
- After the ingress PE router for an NG MVPN instance performs a GRES event, the egress PE routers could fail to install a new forwarding state for the multicast traffic. Clearing the BGP session on the ingress router can restore traffic to all egress routers. [PR/441392]

High Availability

- On a router with dual Routing Engines and nonstop active routing (NSR) enabled, if you perform a commit synchronize operation when the backup Routing Engine is not available, routing protocol sessions might not be reestablished. To expedite protocol synchronization, issue the restart routing command on the backup Routing Engine when it comes up. [PR/277993]
- In a routing matrix, if you include the prefix-action statement at the [edit firewall family inet] hierarchy level and perform an ISSU operation, the FPCs on the T640 routing nodes do not come online. In the output from the show chassis fpc command, the FPC state is reported as ISSU Error. [PR/391266]
- On M-series, MX-series, and T-series routing platforms, if you configure IPv6 on an interface with no MAC address (such as a SONET or loopback interface), it might cause the Routing Engine to restart. As a workaround, do not configure IPv6 addresses on interfaces that do not have MAC addresses. [PR/439252]

Class of Service

- The class-of-service process (cosd) can fail under certain circumstances when container interfaces (for example, rlsq) and graceful routing engine switchover (GRES) are configured. There is no workaround. [PR/466104]
- 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 protocols l2circuit neighbor *address* interface *interface-name*], [edit (logical-routers *logical-router-name* | logical-systems *logical-system-name*) protocols l2circuit neighbor *address* interface *interface-name*], [edit routing-instances *routing-instance-name* protocols l2vpn], or [edit (logical-routers *logical-router-name* | logical-systems *logical-system-name*) 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]
- On M320 and T-series routing platforms, if you map multiple forwarding classes to the same queue (specify the same value for the queue-num statement at the [edit class-of-service forwarding-classes class *class-name*] hierarchy level for multiple classes) and then include the multiple classes in one scheduler map (by including the forwarding-class statement for each one at the [edit class-of-service scheduler-maps *map-name*] hierarchy level), the commit operation fails with the message "Total bandwidth allocation exceeds 100 percent for scheduler-map." [PR/103370]
- 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]
- On MX960 routers, bandwidth sharing across high priority and strict-high priority schedulers might not be as expected. This issue occurs when the schedulers are configured on logical interfaces. [PR/265603]
- When you set the port speed of a multi-rate POS type 2 PIC to OC3, it does not correctly change the CoS speed value within the PFE. The speed is left at OC12.

This will result in unexpected class-of-service (CoS) behavior and there is no workaround at this time. [PR/279617]]

- When a core-facing interface on a PE router that is acting as an IGP peer is deactivated (for example, by deactivating the interface *interface-name* statement at the [edit protocols ospf area *area-id*] hierarchy level), the following message might be written to the system log: "COSMAN: cosman_unbind_update_if_refcount: Failed to find the ifd *interface-name* (*index*) in the ifdtable for ifl *index*." There is no operational impact. [PR/291630]
- When the sum of shaping-rates at the logical interfaces is greater than the interface bandwidth and the rate-limit statement is applied to one of the logical interface queues, the limiting bandwidth for the queue is based on a scaled down logical interface shaping-rate value rather than the configured logical interface shaping-rate. [PR/441413]

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 then sample and then next-term statements. The workaround provides the same functionality as the other configuration, but avoids the problem behavior. [PR/70473]
- On T-series routers, if there is an ingress firewall configured to drop all incoming multicast packets, the discarded multicast packets are incorrectly sent to the Routing Engine. This causes a high utilization of the CPU (50 percent) on the FPC. [PR/239268]
- Do not use the virtual LAN (VLAN) variable when configuring ether-type or vlan-ether-type match conditions for a firewall filter at the [edit firewall *family* vpls filter *filter-name* term *term-name*] hierarchy level. Using the VLAN variable will cause the firewall filter to fail. [PR/273448]
- The show interfaces filters and show interfaces extensive CLI commands do not display the interfaces. [PR/295977]
- Under some circumstances, when you add a prefix at the [edit policy-options prefix-list *list-name*] hierarchy level, the commit operation might fail with one of the following error messages: “Check-out failed for Firewall daemon (/usr/sbin/dfwd) without details" or "configuration check-out failed." [PR/305510]
- The following message might be written to the system log: "rts_cos_get_shaping_rate_for_ifl(): Entry not found for IFL *index* in cos ifl table" under the following conditions:
  - You configure interface-specific input and output filters that contain logical bandwidth policers (include the logical-bandwidth-policer statement at the [edit firewall policer *policer-name*] hierarchy level, and both that policer and the interface-specific statement at the [edit firewall family *family* filter *filter-name* term *term-name* then] hierarchy level).
  - You apply the filters to an interface (include the input *filter-name* and output *filter-name* statements at the [edit interfaces *interface-name* unit *logical-unit-number* family *family* filter] hierarchy level).

- You apply a traffic control profile to the interface (include the *profile-name* statement at the [edit class-of-service traffic-control-profiles] hierarchy level and the output-traffic-control-profile *profile-name* statement at the [edit class-of-service interfaces *interface-name*] hierarchy level).
- The router receives host-bound packets or IP option packets.

As a workaround, include the shaping-rate statement at the [edit class-of-service traffic-control-profiles *profile-name*] hierarchy level. [PR/314292]

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: vrrpMIB, jnxCosIfqStatsTable, jnxCosQstatTable. [PR/63045]
- The TCP dump is reports a max-response-time within IGMP in seconds while displays units of 1/10th of a second. [PR/424618]

J-Web

- While configuring VAP, the No Broadcast SSID is selected even if the user disables it and redisplays the page. [PR/462903]

Resolved Issues

This section lists issues that are fixed in the JUNOS Release 9.3R4. The identifier following the description is the tracking number in our bug database.

Platform and Infrastructure

- If too many statistics requests are sent to the FPC from the Routing Engine, the kernel might run out of buffers and this results in a Routing Engine failure. [PR/281458: This issue has been resolved.]
- On MX, M120 or M320 routers, with E-3FPC platforms a logical interface flap may trigger a jtree memory leak. [PR/403472: This issue has been resolved.]
- A large volume of next-hop changes in a short period may cause a small number of packets to be lost or sent to the wrong destination. [PR/411567: This issue has been resolved.]
- If a duplicate address is detected for theIPv6 family on an Ethernet interface, the DAD does not restart even after the interface goes down and then back comes up. The has been fixed in JUNOS Release 9.3 and later and in JUNOS software released after April 23, 2009. [PR/421241: This issue has been resolved.]
- The input statistics of the AE interface shows the wrong value if the member link is part of an IQ-2 PIC. [PR/429771: This issue has been resolved.]
- On MX-series and M120 routers, and M320 routers with an Enhanced III FPC, if the VRF configuration includes the vrf-table-label statement, a DPC or FPC might dump core when an MPLS packet with time-to-live (TTL) equal to 0 (zero) or 1

(one) is processed at the egress provider edge (PE) router. [PR/436017: This issue has been resolved.]

- In JUNOS Release 9.1 or earlier, when MVPN is configured with auto-RP and there is a change in the RP of the default routing instance, then an RP address changes and a Layer 2 descriptor leak occurs. [PR/436637: This issue has been resolved.]
- SCU configuration causes the PFE to drop some host-bound packets on M320 and T-series routers. [PR/438261] [PR/438261: This issue has been resolved.]
- Under certain circumstances an Intelligent Queuing PIC might not be able to boot properly on an E3-FPC. [PR/438678: This issue has been resolved.]
- When the FPCs for T1600-FPC4-ES, T640-FPC4-1P-ES, T640-FPC1-ES, T640-FPC2-ES, and T640-FPC3-ES receive corrupted cells through high-speed links, they might unnecessarily reboot and report the following system log error message: "Unrecoverable Error: Flist gtop bit toggled !." No reset is needed to recover from this condition. [PR/441844: This issue has been resolved.].
- On T1600, TX Matrix, or T640 routers installed with one of the following Flexible PIC Concentrators (FPCs)—T1600-FPC4-ES, T640-FPC4-1P-ES, T640-FPC4-ES, T640-FPC1-ES, T640-FPC2-ES and T640-FPC3-ES—and JUNOS Release 9.3 or higher, jtree memory might get corrupted once routes are deleted while traffic is send to those prefixes. This can result in permanent or transient packet drops. One or more of the following symptoms might be logged in the system log:
  - SRCHIP(1): 131072 Discards - stack underflow
  - SRCHIP(1): 129735 Discards - truncated key - next hop
  - SRCHIP(1): 4670347 Multicast list discard route entries
  - SRCHIP(1): SOF (58) > = DMA length (46) (Read Channel)
  - SRCHIP(1): RKME int_status 0x300
  - SRCHIP(1): 14486 Discards - illegal BTT
  - SLCHIP(1): 1617082 new errors (illegal link) in DESRD last stream 0 last lout_key 0xabd0e
  - SLCHIP(1): 1622998 new errors (packet error) in HDRF, lout_hdrf_poll_stats

  There is no workaround and an FPC reboot might be needed to recover. [PR/443171: This issue has been resolved.]
- The kernel may have an error due to the loss of a watchdog if several packets are sent out from the Routing Rngine through an aggregated (SONET) interface when the logical interface is down and the physical interface is up. [PR/449361: This issue has been resolved.]
- FIPS 140-2 Level 2 mode operation is not supported, when dual Routing Engines are on the router. [PR/449750: This issue has been resolved.].
- On MX-series tunnel interfaces configured on DPC show traffic incorrectly on other interfaces. [PR/450844: This issue has been resolved.]

- In a Layer 3 VPN PE carrying multicast routes, an error in the kernel crash might occur when an FPC homing on an aggregate Ethernet interface is restarted. [PR/452999: This issue has been resolved.]
- The FPC experiences a heap memory leak when Ethernet OAM protocols are configured. The workaround is to disable the Ethernet OAM protocols. [PR/453842: This issue has been resolved.]
- Due to a JUNOS software issue, an M120 FEB/FPCx can overreact to a CPU Layer 2 cache single-bit-error. [PR/457157: This issue has been resolved.]

User Interface and Configuration

- During commit synchronize, the backup Routing Engine logs the commands to the TACACS+ server. As a result, the commit synchronize process takes a long time to commit. [PR/424255]
- Wildcard apply groups do not work properly in JUNOS Releases 9.1, 9.2, 9.3R1, and 9.3R2. [PR/425355: This issue has been resolved.]
- Issuing the set cli complete-on-space off command may result in unexpected CLI authorization behavior. [PR/426916: This issue has been resolved.]
- SSH/Telnet sessions may time out for a longer period of time then usual if a user or password is not provided. [PR/428116: This issue has been resolved.]
- The idle sync-other-re process may be incorrectly shown in configuration mode. [PR/433164: This issue has been resolved.]
- If you configure the traceoptions statement under system scripts commit, the router may have commit errors. [PR/438289: This issue has been resolved.]

Interfaces and Chassis

- On MX-series routers configured for graceful Routing Engine switchover (GRES), aggregated interfaces might not operate correctly after any of the following events occurs: (a) a simultaneous reboot of both master Routing Engines, (b) a power cycle of the chassis, or (c) a graceful switchover from a master Routing Engine to the backup Routing Engine. To restore functioning, on the master Routing Engine either issue the commit synchronize full command or restart the interface process (dcd). [PR/309716 : This issue has been resolved.]
- When you reboot an FPC while it is coming online and if the FPC adding process is interrupted before it successfully completes, the chassis process does not operate properly. [PR/400676: This issue has been resolved.]
- Incorporating changes to the interfaces configuration results in a small leak in the DCD process. The leak is at the rate of 16 bytes per interface configured per commit. [PR/411596: This issue has been resolved.]
- When you configure LACP on an aggregated Ethernet interface, the counters displayed by the show interface extensive command might show unexpected values. This problem occurs for logical interfaces that have an incoming interface index value that matches the default index of the data stream. [PR/418054: This issue has been resolved]

- The PPP MTU value of an interface protocol on a peer might change as a result of an irrelevant configuration change and cause the PPP MTU negotiation to fail. [PR/421706: This issue has been resolved.]
- When you change a hardware Field Replacement Unit (FRU) in the chassis, the craft process (craftd) might fail upon reinitializing the device list and generate a core file. This does not affect normal operation of the FRU. [PR/429171: This issue has been resolved.]
- On MX480 and MX960 platforms, the FAN LED stays green even when the FAN tray is pulled out. [PR/429521: This issue has been resolved.]
- The algorithm that is responsible to switch over the SFM and take the FPC offline does not clear the errors (hard/soft) on each FPC after the SFM is switched over. [PR/433616: This issue has been resolved.]
- For some interfaces, when configured with the WAN-PHY framing mode, the monitor interface command might be missing some counters. [PR/435775: This issue has been resolved.]
- A large number of ATM2 error interrupts might cause the FPC to fail. [PR/438073: This issue has been resolved.]
- In the output of the show chassis pic fpc-slot x pic-sloty command, the SFP-GE40KM SFP might be shown erroneously as 1000LH instead of 1000EX. [PR/438753: This issue has been resolved.]
- When the same logical interface is deleted from the default system and added into a logical system, the Routing Rngine might fail. [PR/441284: This issue has been resolved.]
- When the sum of shaping rate at a logical interface is greater than the interface's bandwidth and a rate limit is applied to one of the logical interface queues, the bandwidth limit for the queue is based on a scaled-down logical interface shaping rate value rather than the configured logical interface shaping rate. [PR/441413: This issue has been resolved.]
- On M-series routers, BGP sessions flap when any configuration change happens, even an relevant one. As a workaround, make the difference between the configured MRRU and the MTU to be greater than eight. [[PR/442688: This issue has been resolved.]
- When the ingress router re-signals an RSVP session, traffic could egress from a disabled SONET interface that is part of an APS group that is using container interfaces. As a workaround, switch the APS interfaces. [PR/443295: This issue has been resolved.]
- If VRRP tracks a cloned route this is because the cloned route will always be treated as down. The reason this it is always treated as down, is that the unicast cloned routes are not added to the routing table. [PR/446408: This issue has been resolved.]

Services Applications

- A TCP-based stateful firewall flow might remain active after the service interface inactivity timeout expires, even though the corresponding TCP session is already closed. Several iterations of Reset and TCP keepalive messages might be exchanged between the peers before the flow is completely closed. [PR/446960: This issue has been resolved.]

General Routing

- The show helper statistics and clear helper statistics commands are not available on MX-series platforms on or after the following JUNOS releases: 9.3R4, 9.4R4, 9.5R3, and 9.6R2. [PR/445240: This issue has been resolved.]

Routing Protocols

- When more than one external path originates from the same autonomous system (AS), the JUNOS software does not comply with the RFC 5004 path selection algorithm. [PR/392819: This issue has been resolved.]
- Deactivation of routing instances might cause the routing protocol process (rpd) to create a soft assertion failure. [PR/396122: This issue has been resolved.]
- In some cases (for example, after a repeated power-down event), one of the internal database files (/var/db/lmpd-name-id.db) becames corrupt, causing the lmpd system process to fail on commit. As a workaround, delete the file and commit again. [PR/403129: This issue has been resolved.]
- If a multiaccess interface is disabled, it is advertised as a disabled link in the router LSA after the Routing Engine switchover. [PR/418559: This issue has been resolved.]
- If OSPF is in overload mode on the standby Routing Engine but not in overload mode on the master Routing Engine, it may take a long time to install OSPF routes on the standby Routing Engine. [PR/421636: This issue has been resolved.]
- In rare cases, the BPG cleans the data structures correctly when the entire peer group fails and the peer group is deleted. [PR/423060: This issue has been resolved.]
- In a large-scale BGP multipath setup, the BGP multipath calculation uses a large amount of CPU and slows down RPD for a long period of time. [PR/424360: This issue has been resolved.]
- If RIP authentication is turned on, updates may get dropped on sequence number mismatch because they are not processed in the order they are received. [PR/429297: This issue has been resolved.]
- The assert condition is not valid for cases where the PIF is flapped. [PR/429392: This issue has been resolved.]
- Community types are being allocated at random to the members in the community list. As a result, extended communities might be treated as simple and vice versa, which causes failures in the VRF import code. [PR/430728: This issue has been resolved.]

- With non-stop routing enabled for BGP, the master and backup RPD instances will fail to establish and maintain a synchronized state. [PR/434162: This issue has been resolved.]
- If a static route is pointing to a discard configuration, a failure may happen when the router attempts to collect the multicast statistic data. [PR/434298: This issue has been resolved.]
- A Layer 3 VPN BGP using the show bgp neighbor command shows local-id 0.0.0.0 as output when NSR is enabled. [PR/434321: This issue has been resolved.]
- With BGP multipath configured, the BGP trace option flags may not be refreshed after a change in the trace-option flag configuration. [PR/436440: This issue has been resolved.]
- Embedded RP is not created upon receiving a trigger from multicast traffic. Deactivate and activate the configuration to fix the problem. [PR/437893: This issue has been resolved.]
- Embedded RP configurations cause continuous RPD failure if PIM is disabled. [PR/438159: This issue has been resolved.]
- When you use auto-rp, if the rendezvous point (RP) configuration is deactivated and then reactivated on the provider edge (PE) router, the router fails to rediscover the RP announced by the customer edge (CE) router. [PR/438356] [PR/438356: This issue has been resolved.]
- If a RIB is referenced within the FROM clause of a policy statement, the statement might change on each commit. This can lead to route flaps on every commit if the statement is used as the import policy for a RIB group, which in turn is referenced in OSPF. [PR/441557: This issue has been resolved.]
- RPD may fail if a VRF routing instance is reconfigured in a single commit from Draft-Rosen MVPN to Next-Gen MVPN with RSVP-TE inclusive provider tunnels. [PR/442391: This issue has been resolved.]
- When you configure the path-selection always-compare-med statement at the [edit protocols bgp] hierarchy level, BGP multipath might not find all eligible paths. [PR/444629: This issue has been resolved.]
- When BGP NSR is configured with sampling (under forwarding-options sampling), duplicate updates for some prefixes could be sent during a Routing Engine switchover. [PR/458669: This issue has been resolved.]

MPLS Applications

- On M-series and T-series routers, when the MPLS label-switched path (LSP) re-optimizes (or changes path) followed by a signaling failure along that path, then the path change does not occur till the next LSP re-optimization event. [PR/401343: This issue has been resolved.]
- The load-balancing spread is affected when both the primary and the first secondary LSP are out of commission. [PR/422596: This issue has been resolved.]
- The mplsResourceTunnelTable reports bandwidth in bps instead of Kbps. [PR/432716: This issue has been resolved.]

- The MPLS LSP auto-bandwidth adjustment may stop working while RSVP signals for the path; either optimization is initiated or the LSP goes down. [PR/4438157: This issue has been resolved.]
- On a PE router, when an uplink is deactivated, the MPLS LSP BFD session over this link may not switch to other uplinks. [PR/454071: This issue has been resolved.]
- When MPLS traceroute is executed in downstream mapping TLV (TLV 2), the reply packet contains misleading values because of an MPLSOAMD error. [PR/454796: This issue has been resolved.]

VPNs

- Applying configuration changes that remove both static P2MP LSP and a static MVPN provider tunnel group configuration, can result in RPD failure. To avoid this problem, first remove the provider-tunnel configuration, then remove the LSP P2MP configuration. [PR/288456: This issue has been resolved.]
- In Layer 2 CCC scenarios packets where the size is less than 64 bytes, the scenarios packets may be erroneously padded when forwarded through an Ethernet uplink. As a result, the packets size arriving at the remote end will not correspond to those that were originally sent. [PR/420037: This issue has been resolved.]
- If you create new VPLS instances with a provider-tunnel Point-to-Multipoint (P2MP) label-switched path template, the routing protocol daemon (RPD) might restart, creating P2MP LSP paths. [PR/442544: This issue has been resolved.]
- While configuring a Layer 2 VPN routing instance, if the protocol's Layer 2 VPN stanza is not included as part of the routing instance configuration when a commit is performed and instead is added during a later commit, the Layer 2 VPN session associated with this routing instance may not come up. [PR/449494: This issue has been resolved.]

High Availability

- When you issue the show chassis ethernet-switch statistics command on a routing platform with graceful Routing Engine switchover (GRES) enabled, the two Routing Engines might be unable to exchange information for about 2 seconds. [PR/233779: This issue has been resolved.]
- The MIB definitions, jnxPicXDpcCombo10X1GE and jnxPicXQDpcCombo10X1GE for Combo DPC PICs, are missing in the database which causes errors in the chassis process (chassisd) logs. [PR/418469: This issue has been resolved.]
- After an ISSU software upgrade on the MX-series router, you might see a kernel database replication error, an ISSU prepare timeout, and a core dump. These problems might be due to issues with allocated schedulers after the ISSU. This issue is seen only with Gigabit Ethernet Enhanced Queuing IP Services DPCs. [PR/427694: This issue has been resolved.]
- The TX LCC displays an error when ARP entries time out and are added back. This problem occurs with JUNOS Release 9.0 and later (released after August 14, 2007) and in JUNOS Release 8.5R3.3 and 8.5 (released after October 17, 2008). [PR/450698: This issue has been resolved.]

Layer 2 Ethernet Services

- For MX480 routers only, the temperature gap between the MX480 fan speed-up and slow-down has changed from 0 degree Celsius to 5 degree Celsius. Before the change, the fan speeds up to a maximum temperature of 54 Celsius and slows down to 53 Celsius (0 degree gap). After the change, the fan speeds up to a maximum temperature of 56 Celsius and slows down to 49 Celsius (5 degree gap). [PR/394651: This issue has been resolved.]
- When you configure GRES on the MX-series router, the SIB might not initialize if you reboot both Routing Engines simultaneously, or reboot the router with only one Routing Engine installed. [PR/408359: This issue has been resolved.]
- When the router is configured as a DHCP relay agent with the option 82 enabled, it starts dropping packets when the packet size exceeds the maximum size as specified in option 57. [PR/411626: This issue has been resolved.]
- The relay-option-60 configuration, located under the group statement, stops working if something else is changed under the same group statement. [PR/434373: This issue has been resolved.]

High Availability

- An AGRES switchover may cause an FPC failure if the interfaces configuration contains the following statement: sp-x/y/0 { unit 0 { family inet; }. [PR/399152: This issue has been resolved.]
- If static routes are configured under [routing-options], which points to a discarded interface, and if GRES is also configured, then the kernel database may not synchronize with the backup Routing Engine after a GRES switchover is performed. The backup Routing Engine displays a connection error. [PR/399888: This issue has been resolved.]
- When the IPv6 protocol is configured in an IP-IP tunnel and if GRES and NSR are enabled, the backup Routing Engine might display a replication error. [PR/420102: This issue has been resolved.]
- Installing OSPF routes may take a longer then normal period of time, if OSPF is in overload mode on a standby Routing Engine and is not in overload mode on the master Routing Engine (RE). [PR/421636: This issue has been resolved]
- When you use auto-RP and if the rendezvous point (RP) configuration is deactivated and then reactivated on the provider edge (PE) router, the router will fail to rediscover the RP announced by the customer edge (CE) router [PR/438356: This issue has been resolved]
- When you configure the path-selection always-compare-med statement at the [edit protocols bgp] hierarchy level, BGP multipath may not find all eligible paths. [PR/444629: This issue has been resolved]

Class of Service

- The packet drop cannot be brought down to zero. However, with this fix the packet drop should be reduced by nearly half. [PR/429961: This issue has been resolved.]
- On M320 routers, when the Tunnel PIC is on a standard FPC, multicast traffic conforming to Internet draft-rosen-vpn-mcast-08.txt might be subject to incorrect CoS queuing and rewrite. [PR/433142: This issue has been resolved.]
- After the aggregate chassis configuration is deactivated then activated, the classifier might not be properly applied on aggregate interfaces. [PR/442240: This issue has been resolved.]
- After an FPC restart, the classifiers might not be properly applied to the aggregate members if they have LACP configured. This following error message is displayed: Jun 4 12:43:02 sting-re1 fpc0 SLCHIP(0): Unable to fathom what channel used by IFL 68 Jun 4 12:43:02 sting-re1 fpc0 SLCHIP(0): error 1 in setting QoS table 1 for ifl 68 Jun 4 12:43:02 sting-re1 fpc0 COSMAN: lchip write failed, lchip 0 while binding IFL(68) to classifier(1) Jun 4 12:43:02 sting-re1 fpc0 SLCHIP(0): Unable to fathom what channel used by IFL 68 Jun 4 12:43:03 sting-re1 fpc0 SLCHIP(0): error 1 in setting QoS table 1 for ifl 68 Jun 4 12:43:03 sting-re1 fpc0 COSMAN: lchip write failed, lchip 0 while binding IFL(68) to classifier(1)

  The problem is seen on JUNOS Release 9.3, 9.4 releases shipped after 08/15/2008. Deactive and activate CoS to fix the problem. [PR/442418: This issue has been resolved.]
- When an Intelligent Queuing PIC is taken offline and then brought back online, the chassis scheduler map might change to [95,0,0,5]. As a workaround, deactivate the chassis scheduler map before taking the PIC offline and then activate the chassis scheduler map after PIC comes back online. [PR/444543: This issue has been resolved.]
- Tail drops are not seen in the Routing Engine CLI output. [PR/446617: This issue has been resolved.]

Forwarding and Sampling

- Policers cannot be modified after a system upgrade because of a flaw in the parser routine. This error occurs when the current item is deleted and then the parser cannot proceed to the next item. With the fix, the routine in the forwarding process (dwfd) has been modified so that the next item in the object tree is fetched before the current object is parsed. [PR/433418]

Network Management

- When the SNMP has a response that is larger than 9KB, a "Message too long" log is reported but no SNMP get response failure occurs. [PR/389559: This issue has been resolved.]
- When subagents are slow in responding to SNMP queries, the SNMP process continues to buffer the incoming SNMP requests. SNMP memory becomes

exhausted after the buffer increases to a bigger value, which causes the SNMP process to fail. [PR/430106: This issue has been resolved.]

- If the master snmpd restarts in a TX Matrix platform and the SNMP subagent running with an LCC chassisd tries to register MIB objects with the master snmpd, the registration progress fails and results in the snmpd (running at SCC) utilizing large amounts of CPU. [PR/438085: This issue has been resolved.]

Previous Releases

Resolved Issues

9.3R3

This section lists issues that were fixed in JUNOS Release 9.3R3. The identifier following the description is the tracking number in our bug database.

Platform and Infrastructure

- 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: 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 the resolve.conf file does not include a proper working DNS server name, the show ntp associations command output displays the message Can't find host *localhost* with NTP server definitions." Because the DNS server name is not mandatory in the resolve.conf file, the error message is unnecessary. [PR/270915: This issue has been resolved.]
- You might encounter output drops with the 10–Gigabit Ethernet PICs. The output drops occur because the software incorrectly calculates the number of queues for polling statistics in a 10-Gigabit Ethernet PIC, even though it is different from other PICs. [PR/277693: This issue has been resolved.]
- On MX-series routers using Routing Engine-based sampling, when samples are sent from the Packet Forwarding Engine to the Routing Engine over certain interfaces, the interface Input/Output index and next-hop address are set to 0. The following interfaces are affected: ge-x/0/y, ge-x/1/y, xe-x/2/0, and xe-x/3/0. [PR/286089: This issue has been resolved.]
- When an IPv6 BGP peer becomes unreachable, the raw IPv6 packets might be forwarded without the correct Layer 2 encapsulation over an Ethernet connection. [PR/314629: This issue has been resolved.]
- The MX-series Tri-rate DPC does not support MAC accounting and returns the following message: "error: MAC accounting and policing not supported." [PR/387919: This issue has been resolved.]
- On T1600 routing nodes with JUNOS Release 9.3R1 or 9.3R2, if there are interface flaps and routes from 0.0.0.0 to 127.255.255.255 using an indirect next hop,

the following error message might be triggered in the syslog: "JTREE(jt_nh_get_reachable_nh32): Not reachable 0x00000000:0x2d740780 for seg 1 (rt_jtree_build_nh)" and forwarding traffic is impacted. [PR/392876: This issue has been resolved.]

- For aggregated interfaces only, when GRES is enabled and the neighboring server fails, the next hop turns to a hold next hop which waits to be resolved. If the next hop is resolved immediately, the replicated Routing Engine (RE) might panic. [PR/394209: This issue has been resolved.]
- In an MPLS Layer 3 VPN network, the **traceroute** command does not return a valid result (it returns three asterisks [* * *] instead) for the hop between two routers when their configuration includes both of the following features:
  - Per-packet load balancing (the **load-balance per-packet** statement is included at the [**edit policy-options policy-statement** *policy-name* **then**] hierarchy level and that *policy-name* statement is included at the [**edit routing-options forwarding-table**] hierarchy level)
  - Multiple equal-cost paths between the routers (for example, when the **encapsulation frame-relay** statement is included at the [**edit interfaces** *interface-name*] hierarchy level for a SONET/SDH interface and the same address is specified for more than one of its logical interfaces at the [**edit interfaces** *interface-name* **unit** *logical-unit-number* **family** *family* **address**] hierarchy level)

  [PR/396280: This issue has been resolved.]
- When you have configured the **vrf-table-label** statement at the [**edit routing-instances** *routing-instance-name*] hierarchy level for a VRF routing instance, IPv4 and IPv6 MTU error notification is not handled properly. On M320 routers with an incoming FPC as SFPC and an outgoing FPC as FFPC, large IPv6 packets are not being detected and discarded properly. [PR/397334: This issue has been resolved.]
- When the Routing Engine requests numerous statistics that surpass a set boundary, "PFEMAN: Couldn't write..." messages might be logged and DPC failures occur. [PR/398233: This issue has been resolved.]
- When the multicast MAC address is configured on the local PE for a CE device, traffic originating from a remote PE is silently dropped without informing the source that the data did not reach its intended recipient [PR/398698: This issue has been resolved.]
- Prolonged fast interface flaps with thousands of ARP entries might cause the FPC to stop functioning. [PR/399175: This issue has been resolved.]
- On T640 and T1600 routing platforms with the Enhanced Scaling FPC4, errors such as the following might be written to the system log: "*x* new errors (mtu error) in HDRF,lout_hdrf_poll_stats," "Error (code: 30, type:Minor) encountered, cmalarm_passive_alarm_signal," and "1 new errors in SLout OP." There is no operational impact. [PR/399258]
- On egress PE routers, the correct EXP classifier is not applied to label-switched interfaces (LSIs) that are created by including the **vrf-table-label** statement at the [**edit routing-instances** *routing-instance-name*] hierarchy level. [PR/399634: This issue has been resolved.]

- In specific a configuration such as MVPN, restarting RPD causes a small memory leak on the PFE lookup table. [PR/400917: This issue has been resolved.]
- For T640 routing nodes only, when you configure per-packet load balancing, the outgoing traffic is dropped. This issue is exacerbated if you configured two PFE routing instances. [PR/402031: This issue has been resolved.]
- When the ifd channel mode is of type HYBRID, LSI statistics are counted every time ifl_stats are collected for each logical interface. This causes the LSI input counters to be incremented by a multiple of the logical interfaces. [PR/404857: This issue has been resolved.]
- On MX-series routers, when IGMP snooping is enabled in a VPLS instance, a VPLS interface flap causes a DPC to unexpectedly restart. [PR/405136: This issue has been resolved.]
- The traffic class byte is set to 0x00 in the header of some BGP packets sent between interfaces that have IPv6 addresses, instead of the correct setting of 0xc0 (INTERNETCONTROL). [PR/406802: This issue has been resolved.]
- For MX-series routers running with JUNOS Release 9.1R1 or higher, when traffic is sent to the router with the IEEE 802.1p value set to 2 or the source class usage (SCU) configured, the packets are discarded when they reach the PFE. [PR/414491: This issue has been resolved.]
- The show pfe statistics CLI command does not display I-CHIP Ipktwr packet drop counts. [PR/414477: This issue has been resolved.]
- Under rare circumstances, the kernel panics on the TX Matrix LCC or on the SRX-series platform following a Routing Engine switchover or an RDP connection timeout between the LCC and SCC. [PR/416973: This issue has been resolved.]
- For multicast traffic, if the OIF is on an aggregated interface and its member link is on a different PFE (for example, 7/1/0 and 6/1/0), multicast traffic might be lost after the FPC, which has IIF for the multicast, is rebooted. [PR/418583: This issue has been resolved.]
- Initial ARP packets are discarded by the default ARP policer because when a T1600 routing nodes FPC restarts, the current credit is initialized to JT_POL_SR_CURRENT_CREDIT_MAX, which is 0xFFFFF. This has a high negative value in SR, so packets are dropped until it goes down. As a workaround, you can initialize the current credit to max_credit_limit (which is equal to (credit_limit / Rate) * time_credit), approximately equal to TC. [PR/419909: This issue has been resolved.]
- The SNMP remote operations process (rmopd) might fail after configuring a BGP neighbor with a local address. [PR/420504: This issue has been resolved.]
- In JUNOS Release 9.3R1 or higher, on Juniper Networks routers with Type 4 FPCs or T1600 routing nodes, multicast traffic is not counted within the interface statistics counters once class-of-service rewrite rules have been applied to the interface. [PR/420681: This issue has been resolved.]
- On the MX-series router, when you configure MPLS and a tunnel configuration on the same Gigabit Element (GE) DPC, the tunnel interface shows traffic as the sum of the traffic of the other Gigabit Element (GE) interfaces on the DPC. This is a cosmetic issue and does not affect functionality. [PR/422274: This issue has been resolved.]

- When an aggregate bundle fails and the aggregate bundle is part of an Equal Cost Multi-Path (ECMP), there is a short transient window while traffic is re-routed where one or all of the following entries is reported in the message log: - PFE: Detected error nexthop - RCHIP(1): RKME int_status 0x10000000 - LCHIP(1): 3067 new errors (illegal size) in DESRD - LCHIP(1): 3067 new errors (illegal link) in DESRD - RCHIP(1): SOF (61) > = DMA length (46). [PR/424741: This issue has been resolved.]
- On MX-series routers, the FPC might reboot without a failure if the DWDM is incorrectly configured. Either disconnect the offending link or configure the Disable statement at the [edit interfaces] hierarchy level to stop the FPC reboots. [PR/430703: This issue has been resolved.]
- When configuring Proxy ARP on unnumbered interfaces, the router can incorrectly answer address collision detection ARP requests, causing DHCP clients to decline the offered address. [PR/431192: This issue has been resolved.]
- When you configure flow monitoring on a T1600 with a T640 or T1600 Enhanced Scaled FPC4 and the input and output traffic are located on the same bottom PFE1, then the next-hop address and output interface are set to 0. [PR/431567: This issue has been resolved.]
- On MX-series, M120, and M320 routers with an Enhanced III FPC, the DPC FPC fails if the VRF configuration includes the vrf-table-label statement when an MPLS packet with time-to-live (TTL) is set equal to 0 (zero) or 1 (one) and is processed at the egress PE. [PR/436017: This issue has been resolved.]
- An ARP retry count is incorrect in that instead of sending out the first five retries every second, the third and consequent retries are sent every 15 seconds. [PR/436580: This issue has been resolved.]
- On MX-series routers with a Combo DPC (20-port 1-Gigabit Ethernet 2-port 10-Gigabit Ethernet), if the family mpls statement is included at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level for any 1-Gigabit Ethernet port of a DPC slot, the show interfaces statistics command reports zero values for input traffic at all ports. This issue does not affect the input traffic statistics for the 10-Gigabit Ethernet ports. This is a cosmetic issue and does not affect functionality. [PR/436653: This issue has been resolved.]

User Interface and Configuration

- The alarm process (alarmd) updates /var/db/feature.db, a license-tracking file, every 60 seconds, even on routers that do not support the JUNOS software licensing feature (for example, the M7i, M10i, M40e, and T-series routing platforms) and causes unnecessary hard disk drive activity. [PR/308466: This issue has been resolved.]
- The container value is unavailable when the commit script show configuration system scripts commit is used with traceoptions and when the direct-access statement is set. [PR/394243: This issue has been resolved.]
- The algorithm that switches over the SFM and takes the FPC offline, does not clear the hard/soft errors on each FPC once the SFM is switched over. [PR/433616: This issue has been resolved.]

- When the direct-access statement is configured, the firewall filter input-list in a commit script may not return an expected value. [PR/406663: This issue has been resolved.]
- The RPC get-configuration statement may not get the expected output if both direct-access and filter are configured under [system scripts commit]. [PR/406687: This issue has been resolved.]
- You get a commit fail when applying a group to the chassis section of a configuration. [PR/425355: This issue has been resolved.]
- When you use the commit confirmed command on TX-series routers, it fails to roll back the original configuration as expected. [PR/425642: This issue has been resolved.]
- If you configure the traceoptions statement under system scripts commit, the router may have commit errors. [PR/438289: This issue has been resolved.]

Interfaces and Chassis

- In the output from the show interfaces extensive command, the count of REI-P errors in the SONET path section is incorrect when the RDI-P error also appears in the SONET defects field. [PR/256049: This issue has been resolved.]
- On aggregated Ethernet interfaces configured for LACP (the lacp statement is included at the [edit interfaces ae*x* aggregated-ether-options] hierarchy level), if you deactivate one of the interfaces in the aggregate, multicast traffic might not be detoured as expected. [PR/313617: This issue has been resolved.]
- On a router with dual Routing Engines, if the hard disk is inoperable or missing on the backup Routing Engine, no chassis alarm is set (visible in the output of the show chassis alarms command), nor is an SNMP trap or system log message generated. The only indication is a line like the following in the output from the show system boot-messages command: "ad*x*: not attached, missing in Boot List." [PR/392837: This issue has been resolved.]
- On the T1/E1 Circuit Emulation PIC, if you specify an invalid value for the payload-size statement at the [edit interfaces (t1 | e1)-*fpc/pic/port* satop-options] hierarchy level, the DS1 alarm LOF is raised, as reported in the output from the show interfaces (t1 | e1)-*fpc/pic/port*:*channel* command.

  The valid values for the payload-size statement are as follows:

  - In T1 mode, a multiple of 24 in the range 24 to 1024
  - In E1 mode, a multiple of 32 in the range 64 to 1024

  [PR/395143: This issue has been resolved.]
- In JUNOS Release 9.3R1 and later, SONET Automatic Protection Switching (APS) does not work correctly on the 4-port Channelized OC3/STM1 Circuit Emulation PIC with SFP. [PR/402068: This issue has been resolved.]
- On channelized OC12 intelligent queuing (IQ) interfaces, incoming code violation path (CV-P) messages might not trigger the sending of remote error indication path (REI-P) messages.[PR/47188: This issue has been resolved.]

- While bringing a PIC online, after bringing a router online and performing an FPC or PIC (re)start, the interface hold-down up timer is activated and the interface comes up immediately. [PR/277236: This issue has been resolved.]
- In TX Matrix platforms, the show chassis fpc X command returns an error instead of showing the FPC information when X is greater than 8. [PR/387950: This issue has been resolved.]
- In OC768-over-OC192 mode on the 4-port OC192c PIC, when you change the clocking internal statement to clocking external at the [edit interfaces *interface-name*] hierarchy level, the clock may not come up. [PR/395847: This issue has been resolved.]
- When the no-auto-negotiation statement is configured under a port within IQ2 PICs, the down link may flap. [PR/397491: This issue has been resolved.]
- On T640 router nodes when the FPC is taken offline, the AE bundle statistics (which issue the monitor interface traffic command) display a high value. This is not an issue for the TX Matrix platform. [PR/399451: This issue has been resolved.]
- Running OAM under an aggregate interface might not detect a link failure in a child interface. This causes the router to direct network traffic to a destination where it is lost. [PR/399868: This issue has been resolved.]
- The output for queue counters under the show interfaces command (xe-fpc/pic/port extensive) might be incorrect when traffic is passed at near maximum throughput to any Queuing IQ2 or IQ2E PICs or DPCs. [PR/401431: This issue has been resolved.]
- High priority traffic gets RED dropped even though the rate is lower than the shaping-rate under the following conditions:
  - 248 VLANs are configured on a single port within an IQ2 Gigabit Element (GE) PIC
  - The shaping rate for each VLAN is set to 4m and a buffer size for high priority traffic (for example, real-time is 5 percent)
  - Both high and low priority traffic are sent out through all 248 VLANs where the total rate is higher than the line rate

  [PR/401893: This issue has been resolved.]
- When Multilink Frame Relay encapsulation is configured on an interface using the encapsulation multilink-frame-relay-uni-nni statement is included at the [edit interfaces *interface-name*] hierarchy level), the kernel might generate an error. [PR/407608: This issue has been resolved.]
- When a 10-Gigabit Element interface of a DPC is connected to a faulty optical card which causes the link state to change at a very high rate, the DPC fails. [PR/411072: This issue has been resolved.]
- When a Layer 2 policer is applied to the egress interface of a router, the dropped frame statistics might show incorrect information. [PR/419181: This issue has been resolved.]
- On an IQ2 PIC, the slow aging interval might be overwritten with a value of 202 seconds which causes the MAC entry to be removed in 6 to 7 minutes. [PR/419510: This issue has been resolved.]

- The address family of child next hops is incorrectly set to the address family of the IFF, instead of the address family of the parent next hop. [PR/425802: This issue has been resolved.]
- A NULL pointer reference in an ifinfo failure is caused by a loss of synchronization with GRES-enabled Routing Engines. [PR/43112: This issue has been resolved.]
- The SFP-GE40KM SFP may display as 1000LH instead of 1000EX in the output of the chassis pic fpc-slot x pic-slot y command. [PR/433616: This issue has been resolved.]

Services Applications

- The issue occurs when you configure the NAT match-direction output statement and attach it to a interface-style service set on an egress interface. When you explicitly configure forward and backward rules for a NAT service set, an ICMP fragmentation-needed message is not sent and the traffic is dropped without notification. If the backward rule is not configured and is left implicit, this problem is not seen. An explicit backward rule causes the ICMP error packet to be handled as a new flow. [PR/238215: This issue has been resolved.]
- On an M7i or M10i router with the enhanced CFEB, if you issue the deactivate forwarding-options sampling command, sampling stops for both IPv4 and IPv6 traffic. If you then issue the activate forwarding-options sampling command, sampling resumes for only IPv4 traffic. [PR/415140: This issue has been resolved.]

General Routing

- When you configure multiple addresses for the from neighbor statement inside a routing policy term, only the last address takes effect. [PR/414768: This issue has been resolved.]
- On TX Matrix platforms, use of generate in the routing-options stanza with reference to a policy results in the commit not completing successfully. [PR/416380: This issue has been resolved.]
- A RPD error occurs after you commit changes to a routing instance configuration. [PR/425126: This issue has been resolved.]

Routing Protocols

- On a router with dual Routing Engines and NSR configured, the backup RPD may go down in rare instances while processing an indirect next-hop delete. [PR/302731: This issue has been resolved.]
- Inefficient deletions of BGP routes from the routing instance table cause the scheduler to slip. [PR/305027: This issue has been resolved.]
- When more than one external path originates from the same autonomous system (AS), the JUNOS software does not comply with the RFC 5004 path selection algorithm. [PR/392819: This issue has been resolved.]
- The deactivation of a routing instance causes an RPD to create a soft assertion failure. [PR/396122: This issue has been resolved.]

- On a router configured with nonstop routing (NSR), when you apply the BGP import policy and then issue the clear bgp neighbor address soft command to reset BGP, the policy does not take effect. [PR/396291: This issue has been resolved.]
- If you specify an IPv6 address as a value in the ssm-groups statement at the [edit routing-options multicast] hierarchy level, the SSM group does not work as expected. [PR/399352: This issue has been resolved.]
- When you enable distributed periodic packet management by including the delegate-processing statement at the [edit routing-options ppm] hierarchy level, BFD packets are transmitted on a queue other than queue 3 (this could be queue 0 or queue 4 depending on the JUNOS software version). [PR/400907: This issue has been resolved.]
- The SNMP interface index is not set internally which causes the MIB interface queries to display the index value as zero.This can also cause SNMP interface MIB queries for statistics to return stale information. [PR/401038: This issue has been resolved.]
- If GRES is not enabled on a Routing Engine switchover, then the routing protocol process (rpd) on the new backup Routing Engine quits before cleaning up the forwarding table. [PR/402372: This issue has been resolved.]
- When you issue the mtrace *source* command and the route to the source is defined in the routing table for a PIM nonforwarding instance (that is, not in the main instance table, inet.0), the command fails with the following messages: "...giving up" and "Timed out receiving responses." [PR/403033: This issue has been resolved.]
- When an operator executes the show route aspath-regex command and then attempts to escape with CTRL + C, an RPD generates a failure. [PR/403410: This issue has been resolved.]
- When the OSPF overload timeout is set, even as low as the minimum of 60 seconds, the external LSA may not be generated even after the overload timer times out. [PR/404097: This issue has been resolved.]
- When peers in different BGP peer groups have similar export policies such that identical advertisements are sent, the routing protocol process (rpd) might generate an error and become unresponsive when the backup Routing Engine comes online. [PR/404471: This issue has been resolved.]
- When certain statements are included at the [edit protocols bgp group *group-name*] hierarchy level, the routing protocol process (rpd) might generate an error and stop operating in some circumstances. [PR/404667: This issue has been resolved.]
- Aggregate routes with a large number of contributing members cause the routing protocol process (rpd) to monopolize the CPU constantly with frequent routing changes. However this condition applies only when you configure a policy with the aggregate-contributor match condition. [PR/405499: This issue has been resolved.]
- An SNMP MIB walk of the downstream interfaces of point-to-multipoint multicast routes might cause the routing protocol process (rpd) to fail. [PR/405505: This issue has been resolved.]
- When rapid configuration commits occur for a certain type of configuration changes that include nonstop routing configuration, rpd may stop consuming

further configuration changes with the message "SIGHUP while previous commit isn't yet complete." [PR/405761: This issue has been resolved.]

- If you redistribute a default route or other labeled unicast FEC with the discard or reject action into BGP and enable traffic statistics at the [edit protocols bgp family inet labeled-unicast] hierarchy level, the routing protocol process (rpd) might fail and FECs might be logged with a value of 0. [PR/407546: This issue has been resolved.]
- When changing from static OSPF and ISIS route load balancing to BGP load balancing with multipath enabled, the routes may not be load balanced correctly until the BPG session is restarted. [PR/407925: This issue has been resolved.]
- PIM mistakenly prefers a specific hidden route over an active less specific route as the RPF route to the MCAST source. [PR/411385: This issue has been resolved.]
- If a multiaccess interface is disabled, it is advertised as a disabled link in the router LSA after the Routing Engine (RE) switchover. [PR/418559: This issue has been resolved.]
- In rare cases, the BPG cleans the data structures correctly when the entire peer group fails and the peer group is deleted. [PR/423060: This issue has been resolved.]
- In a large-scale BGP multipath setup, the BGP multipath calculation uses a large amount of CPU and slows down RPD for a long period of time. [PR/424360: This issue has been resolved.]
- If RIP authentication is turned on, updates may get dropped on sequence number mismatch because they are not processed in the order they are received. [PR/429297: This issue has been resolved.]
- The assert condition is not valid for cases where the PIF is flapped. [PR/429392: This issue has been resolved.]
- Community types are being allocated at random to the members in the community list. As a result, extended communities might be treated as simple and vice versa, which causes failures in the VRF import code. [PR/430728: This issue has been resolved.]
- If a static route is pointing to a discard configuration, a failure may happen when the router attempts to collect the multicast statistic data. [PR/434298: This issue has been resolved.]
- A Layer 3 VPN BGP using the show bgp neighbor command shows local-id 0.0.0.0 as output when NSR is enabled. [PR/434321: This issue has been resolved.]
- With BGP multipath configured, the BGP trace option flags may not be refreshed after a change in the trace option flag configuration. [PR/436440: This issue has been resolved.]
- Embedded RP configurations cause continuous RPD failure if PIM is disabled. [PR/438159: This issue has been resolved.]
- On a router configured for NSR, when you apply a BGP import policy and issue the clear bgp neighbor *address* soft command to reset BGP, the policy does not take effect. (In terms of configuration statements, the nonstop-routing statement is included at the [edit routing-options] hierarchy level and the import *policy-name* statement at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.) As a workaround, either disable NSR or issue the clear bgp neighbor *address*

command without the soft option, which forces BGP peers to reestablish their sessions. [PR/396291: This issue has been resolved.]

- When two BGP peers establish a session, they negotiate the hold time to use for keepalive messages. If one of the peers uses a nondefault hold-time value (that is, the hold-time statement is included at the [edit protocols bgp group *group-name*] hierarchy level in its configuration), and either of the peers goes down immediately after the session is established, the hold timer incorrectly expires after the default interval instead of the negotiated interval. [PR/396823: This issue has been resolved.]
- If you specify an IPv6 address as a value for the ssm-groups statement at the [edit routing-options multicast] hierarchy level, the SSM group does not work as expected. As a workaround, specify only IPv4 addresses. [PR/399352: This issue has been resolved.]
- When you enable distributed periodic packet management (by including the delegate-processing statement at the [edit routing-options ppm] hierarchy level), BFD packets are transmitted on a queue other than queue 3 (queue 0 or 4 depending on the JUNOS version). If system load allows it, disable distributed PPM as a workaround. [PR/400907: This issue has been resolved.]
- When you issue the mtrace *source* command and the route to the source is defined in the routing table for a PIM nonforwarding instance (that is, not in the main instance table, inet.0), the command fails with the following messages: "...giving up" and "Timed out receiving responses." [PR/403033: This issue has been resolved.]
- When certain statements are included at the [edit protocols bgp group *group-name*] hierarchy level, the routing protocols process (rpd) might generate an error and stop operating in some circumstances. [PR/404667: This issue has been resolved.]

MPLS Applications

- After a link flap which triggers a print-to-multipoint LSP reroute, the CCC connection stays down for an long period of time due to a race condition between CSPF runs and the RSVP. [PR/280259: This issue has been resolved.]
- Traffic loss might occur during an LSP switchover. [PR/392406: This issue has been resolved.]
- When you change the configuration of a secondary (standby) LSP in certain ways, the entire LSP is taken down and set up again, which might cause traffic loss or delay. Specifically, the problem occurs if you add or change the value of certain statements at the [edit protocols mpls label-switched-path *lsp-name* secondary *lsp-name*] hierarchy level, including admin-group, hop-limit, and priority. [PR/394184: This issue has been resolved.]
- On M-series and T-series routers, when the MPLS label-switched path (LSP) re-optimizes (or changes path) followed by a signaling failure along that path, then the path change does not occur till the next LSP re-optimization event. [PR/401343: This issue has been resolved.]
- If an RSVP LSP configured with LDP tunnel initiates auto-bandwidth adjustment, the LDP might fail to send keepalive message. This can trigger an LDP session flap as a result of hold-down timer expiration. [PR/407707: This issue has been resolved.]

VPNs

- The time-to-live (TTL) threshold value is not propagated correctly for VPNs that use IPv6 addresses. This might cause multiple entries for the same address in the output from the traceroute command. [PR/257497: This issue has been resolved.]
- When you reboot a PIC or FPC that houses a virtual tunnel (vt-) interface, the interface is not re-created. As a workaround, deactivate and reactivate the interface in the configuration. [PR/266170: This issue has been resolved.]
- When deleting a Layer 2 VPN routing instance and then adding a new VPLS routing instance using the same interface within the same commit, RPD fails. [PR/291407: This issue has been resolved.]
- If you take a PIC offline that hosts a large number (for example, 1000) of CE-facing interfaces in a Layer 2 VPN, the routing protocols process (rpd) might generate an error. [PR/300601: This issue has been resolved.]
- On a router configured for nonstop routing (NSR), if you perform the following sequence of steps, the routing protocol process (rpd) on the backup Routing Engine might generate a failure:

  1. Remove a Layer 2 VPN routing instance (that is, one for which the configuration includes the instance-type l2vpn statement at the [edit routing-instances *routing-instance-name*] hierarchy level).
  2. Commit the configuration.
  3. Immediately create a new Layer 2 VPN routing instance.
  4. Commit the configuration.

  [PR/401057: This issue has been resolved.]
- In a VPLS dual-homed configuration, traffic loss might occur for approximately 20 seconds during a switchover from the backup to the primary interface. [PR/404605: This issue has been resolved.]
- On a router configured as a Layer 2 VPN ASBR or route reflector, if a BGP session to a Layer 2 VPN peer (Layer 2 VPN signaling is enabled) flaps or is explicitly cleared, the backup routing protocol process (rpd) might fail and restart. [PR/407820: This issue has been resolved.]
- If MAC addresses are learned within a VPLS instance, CE devices will communicate directly even though no-local-switching is configured. [PR/419976: This issue has been resolved.]
- Multicast group addresses ending with .232 are classified as SSM groups when using multicast VPNs. These routes are not installed in a multicast VPN routing table and all traffic to these destinations is dropped. [PR/426811: This issue has been resolved.]
- While handling the ifl mismatch notification, multicast code finds the active route from the route (S,G) that should get installed in the forwarding plane which leads to a mismatch. The multicast code then hands the mismatch notification to the protocol that owns the active route. While finding the active route, multicast

ignores the MVPN route and the mismatch notification is dropped. [PR/431211: This issue has been resolved.]

Layer 2 Ethernet Services

- For MX480 router only, the temperature gap between the MX480 fan speed-up and slow-down has changed from 0 degree Celsius to 5 degree Celsius. Before the change, the fan speeds up to a maximum temperature of 54 Celsius and slows down to 53 Celsius (0 degree gap). After the change, the fan speeds up to a maximum temperature of 56 Celsius and slows down to 49 Celsius (5 degree gap). [PR/394651: This issue has been resolved.]
- When you configure GRES on the MX-series router, the SIB might not initialize if you reboot both Routing Engines simultaneously, or reboot the router with only one Routing Engine installed. [PR/408359: This issue has been resolved.]
- When the router is configured as a DHCP relay agent with the option 82 enabled, it starts dropping packets when the packet size exceeds the maximum size as specified in option 57. [PR/411626: This issue has been resolved.]
- The relay-option-60 configuration, located under the group statement, stops working if something else is changed under the same group statement. [PR/434373: This issue has been resolved.]

High Availability

- An AGRES switchover may cause an FPC failure if the interfaces configuration contains the following statement: sp-x/y/0 { unit 0 { family inet; }. [PR/399152: This issue has been resolved.]
- If static routes are configured under [routing-options] which points to a discarded interface and if GRES is also configured then the kernel database may not synchronize with the backup Routing Engine (RE) after a GRES switchover is performed. The backup Routing Engine (RE) displays a connection error. [PR/399888: This issue has been resolved.]
- When the IPv6 protocol is configured in an IP IP tunnel and if GRES and NSR are enabled, the backup Routing Engine (RE) might display a replication error. [PR/420102: This issue has been resolved.]

Class of Service

- When you use wildcards to configure class-of-service (CoS) attributes for interfaces on intelligent queuing PICs (for example, IQ and IQ2), the scheduler map specified for the interface can be applied to the chassis stream. Performing a Routing Engine (RE) switchover in this condition can result in the chassis scheduler map being removed. As a workaround, you can explicitly configure a chassis scheduler map with the scheduler-map-chassis statement at the [edit class-of-services interfaces] hierarchy level. [PR/425710: This issue has been resolved.]
- When you apply a class-of-service (CoS) classifier to a logical interface that has the * (wildcard) value configured as the unit number, the classifier is removed after a Routing Engine reboot occurs. This issue does not occur if the logical interface unit value is configured as a specific numerical value. To apply a CoS

classifier to a logical interface, include the classifier *classifier-name* statement at the [edit interfaces *interface-name* unit unit-value] hierarchy level. For classifier-name, include the name of a classifier configured at the [edit class-of-service classifiers] hierarchy level. [PR/427848: This issue has been resolved.]

- In JUNOS Release 8.4 and later, the commit or commit-check operation fails if a rewrite rule is defined at both 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: This issue has been resolved.]

Forwarding and Sampling

- A flow route is assigned an internal identifier that captures the values of all match conditions specified at the [edit routing-options flow route *route-name* match] hierarchy level. If the length of the identifier exceeds a certain limit, the MIB II process (mib2d) might repeatedly generate an error and fail to restart. The higher the number of match conditions, and the more values specified for conditions that accept multiple values (such as the destination-port and source-port statements), the more likely the problem is to occur. As a workaround, limit the number of conditions or values or both. [PR/273373: This issue has been resolved.]
- On an MX480 router, if you change link speed for a physical interface (by changing the value of the speed statement at the [edit interfaces *interface-name*] hierarchy level) and a rate-limiting output policer is applied to one of its logical interfaces (the output statement is included at the [edit interfaces *interface-name* unit *logical-unit-number* family *family-name* policer] hierarchy level), the traffic rate does not change (as reported by the show policer command). As a workaround, deactivate the policer statement, commit, reactivate the statement, and commit again. [PR/314143: This issue has been resolved.]
- For a filter whose last term has a next-term statement, if the filter 1) is applied individually, and 2) is within the term of another filter, or is applied in an input-list or an output-list, then the firewall process will commit with errors in the log and the filters might not be applied. [PR/395561: This issue has been resolved.]
- The password statement configured under [accounting-options file x archive-sitesis] may not work correctly. [PR/396648: This issue has been resolved.]
- A sample core error occurs when you perform an assertion which causes a memory allocation failure. [PR/418126: This issue has been resolved.]
- When a filter term has "next term" as the action, the action may be shown in the firewall log as "unknown" for the matched outgoing packets. [PR/421810: This issue has been resolved.]
- For list filters, the firewall compiler (dfwc) creates temporary interface-specific filters marked with a flag (DFW_FLAGS_IFACE_INLINE) and uses them to clone as needed. These filters are usually purged from the system after cloning, but with this issue the filters are not purged and occupy index space. The workaround is to identify the unpurged list filter by checking for the flag DFW_FLAGS_IFACE_INLINE and then deleting it manually. [PR/426137: This issue has been resolved.]

- The commit fails with the error message “Referenced prefix-list *xxx*” is not defined under the following conditions:
  - An input-list or output-list is configured on an interface in a logical system
  - The filters in the list are defined under the firewall hierarchy of the main router
  - A prefix list defined under the policy-options of the main router is referenced by one of the filters in the list

  [PR/427253: This issue has been resolved.]
- Policers could not be modified after a system upgrade because of a flaw in the parser routine. This error occurs when the current item is deleted and then the parser cannot proceed to the next item. With the fix, the routine in the forwarding process (dwfd) has been modified so that the next item in the object tree is fetched before the current object is parsed. [PR/433418]

General Routing

- The jnxFWCounterPacketCount MIB module does not show the correct values and displays a zero even if the statistics used in the show command are non-zero. [PR/403563: This issue has been resolved.]
- If the kernel is slow to respond to interface statistics requests made by the Management Information Base II (MIB II) process (mib2d), it could be that the MIB II process is blocking the request. In addition, if there is an interface flap (link down followed by up), the MIB II process may recognize only the latest interface link state and thereby miss modifying the ifLastChange object identifier (OID) associated with the interface, and also miss sending a link down trap. [PR/421585: This issue has been resolved.]
- The Management Information Base II (MIB II) process (mib2d) core is generated when the Routing Engine 1 (RE1) is reloaded. [PR/436218: This issue has been resolved.]

9.3R2

This section lists issues that fixed in JUNOS Release 9.3 R2. The identifier following the description is the tracking number in our bug database

Platform and Infrastructure

- When the Routing Engine hard disk fails, the compact flash might be removed from the list of media used at boot time, instead of the hard disk being removed. In some cases, this makes the Routing Engine unable to initialize. [PR/389540: This issue has been resolved.]
- On M120 and MX-series routers, and on some FPCs on M320 routers, the Packet Forwarding Engine might not free memory correctly during operations on multicast next hops. [PR/396903: This issue has been resolved.]
- On a T1600 routing node, an FPC might stop operating while processing an ICMP TTL expiration packet. Such packets increment the count in the ttl expired field

of the output from the show pfe statistics ip icmp command. [PR/398059: This issue has been resolved.]

- On egress PE routers, the correct EXP classifier is not applied to label-switched interfaces (LSIs) that are created by including the vrf-table-label statement at the [edit routing-instances *routing-instance-name*] hierarchy level. [PR/399634: This issue has been resolved.]
- When you install an FPC in all eight slots on a T1600 routing node configured for graceful Routing Engine switchover (the graceful-switchover statement is included at the [edit chassis redundancy] hierarchy level), the routing node might reboot repeatedly. As a workaround, disable GRES or remove one FPC. [PR/400267: This issue has been resolved.]

User Interface and Configuration

- When you issue the request system (halt | power-off | reboot) other-routing-engine lcc *routing-node-index* command on a TX Matrix platform, the requested operation is performed on the TX Matrix platform instead of the specified routing node (line-card chassis, or LCC). As a workaround, issue the command on the routing node itself (without the lcc option). [PR/241274: This issue has been resolved.]
- On routers that do not use JUNOS software licensing (for example, the M7i, M10i, M40e, and T-series routing platforms) the alarm process (alarmd) nevertheless updates a license-tracking file every 60 seconds. This causes excessive disk activity. As a workaround, become the root user and create an empty directory called /config/license. To determine if a router supports licensing, issue the show system license command. On routers that do not support licensing, the command returns the message "syntax error, expecting < command > " and we recommend the workaround. [PR/308466: This issue has been resolved.]

Interfaces and Chassis

- On MX-series routers, when a DPC configured with a large number of interfaces restarts, the chassis process (chassisd) might write the following messages to the log: "failed to complete channel bonding" and "reached link 5 max index value." [PR/292057: This issue has been resolved.]
- When only one Routing Engine is installed in an M120 router, on the craft interface the LEDs for the power supplies never light up. Similarly, in the PS LEDs section of the output from the show chassis craft-interface command, there is a period in all four fields (indicating that no LEDs are lit). [PR/302504: This issue has been resolved.]
- When Multilink Frame Relay encapsulation is configured on an interface (the encapsulation multilink-frame-relay-uni-nni statement is included at the [edit interfaces *interface-name*] hierarchy level), the kernel might generate an error. [PR/408066: This issue has been resolved.]

Services Applications

- Network address translation (NAT) is not performed correctly for Real-Time Streaming Protocol (RTSP) methods when the Content-Length field is set to 0 (zero). [PR/393171: This issue has been resolved.]

Subscriber Access Management

- If you create multiple subscriber sessions on a logical interface at the same time, some clients might not initialize correctly. The show dhcp server binding detail command reports the value act-prof in the State column for these clients. [PR/303778: This issue has been resolved.]

Layer 2 Ethernet Services

- When more than one of a physical interface's logical interfaces is associated with a bridge domain (the family bridge statement is included at more than one [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level and each logical interface is specified as the value for an interface *interface-name* statement at an [edit bridge-domains *domain-name*] hierarchy level), the monitor *physical-interface-name* command displays incorrect values in the Input packets field of the Traffic statistics section. [PR/397745: This issue has been resolved.]

Routing Protocols

- On a router with dual Routing Engines that is configured for nonstop active routing (NSR) and graceful Routing Engine switchover, if the backup-router or inet6-backup-router statement is included at the [edit system] hierarchy level, the static route to the backup destination is not deleted on the backup Routing Engine when you activate NSR. [PR/305597: This issue has been resolved.]
- If the route to a multicast source address is learned using BGP and the upstream interface goes down, PIM might not detect the outage. As a consequence, the value unknown appears in the Upstream interface and Upstream neighbor fields of the output from the show pim join extensive command. [PR/397410: This issue has been resolved.]
- If PIM sources are accessed via different addresses on the same neighbor, and PIM is deactivated and reactivated on the neighbor, the Upstream interface and Upstream neighbor fields of the output from the show pim join extensive command continue to report the value unknown after the neighbor is active. [PR/400573: This issue has been resolved.]
- When peers in different BGP peer groups have similar export policies such that identical advertisements are sent, the routing protocols process (rpd) might generate an error and become unresponsive when the backup Routing Engine comes online. [PR/404471: This issue has been resolved.]

MPLS Applications

- When the load-balance bandwidth statement is included at the [edit protocols rsvp] hierarchy level on a router with two LSPs to a destination, the balance coefficient is set to zero for the next-hop interfaces in the MPLS forwarding table entry for the route to the destination that is marked with (S=0) (in other words, in the output from the show route forwarding-table family mpls extensive command, the record with the header Destination: *index*(S=0) has Next-hop interface entries where the Balance field does not appear). [PR/257570: This issue has been resolved.]

  When both CSPF and link protection are enabled, in rare instances the routing protocol process (rpd) might generate an error and restart. [PR/266126: This issue has been resolved.]

High Availability

- On an MX-series router configured for VRRP for IPv6, during a mastership change the original master does not relinquish mastership, with the result that both it and the original backup are reported as master in the VR state field of the output from the show vrrp summary command. [PR/398399: This issue has been resolved.]
- On a router configured for nonstop active routing (NSR), if you perform the following sequence of steps, the routing protocols process (rpd) on the backup Routing Engine might generate an error: remove a Layer 2 VPN routing instance (that is, one for which the configuration includes the instance-type l2vpn statement at the [edit routing-instances *routing-instance-name*] hierarchy level), commit the configuration, immediately create a new Layer 2 VPN routing instance, and commit the configuration. [PR/401057: This issue has been resolved.]

Class of Service

- When you update a CoS rewrite rule, the changes are not applied to active multicast streams, but only to streams created after the change. As a workaround, clear all active multicast streams after updating the rule. [PR/266341: This issue has been resolved.]

9.3R1

This section lists issues that were fixed in JUNOS Release 9.2R1. The identifier following the description is the tracking number in our bug database.

Platform and Infrastructure

- 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.]
- 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/238960: This issue has been resolved.]
- When you issue the file copy command with an FTP path as the source or destination and include the source-address option, the specified source address is not used for establishing a connection with the peer FTP server. [PR/240580: This issue has been resolved.]
- On MX960 routers, if you issue the request system power-off other-routing-engine command to power down a Routing Engine, it does not power back on when you then issue the request system power-on other-routing-engine command. [PR/253061: This issue has been resolved.]
- When you configure aggregated interfaces as core-facing links, translational cross-connect (TCC) might not work properly. [PR/267867: This issue has been resolved.]
- Including the mirror-flash-on-disk statement at the [edit system] hierarchy level has no effect. [PR/268474: This issue has been resolved.]
- On MX-series Ethernet Services routers, if the label-switched interface (LSI) is enabled for an xe member link that is part of an aggregated Ethernet (ae) interface, the xe interface statistics are counted twice. [PR/274396: This issue has been resolved.]
- 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: This issue has been resolved.]
- On a Routing Engine of type RE-3.0 (as reported by the show chassis hardware command) with a 1-GB compact flash card, issuing the request system snapshot command might corrupt one or more JUNOS package files in the /altroot/packages directory. [PR/291295: This issue has been resolved.]
- In an environment with many active multicast routes and one or more aggregated interfaces as downstream interface, when an aggregated interfaces flaps or an FPC containing an aggregated interface restarts, the kernel might restart unexpectedly. This issue is seen in networks with greater than 1000 multicast routes. The chance of kernel restarts increases as the number of multicast routes increases or the number of downstream aggregated interfaces increases. [PR/292521: 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/293212: This issue has been resolved.]

- When a Multilink Point-to-Point Protocol (MLPPP) link is incorrectly added to a Multilink Frame Relay (MLFR) bundle, the kernel resets unexpectedly. [PR/294885: This issue has been resolved.]
- An MPLS frame with an explicit NULL label designated for the Routing Engine might be dropped by the Packet Forwarding Engine. [PR/298967: This issue has been resolved.]
- For individual T1 links in an MLPPP bundle, the counts of input bytes and input packets are not reported correctly in the Traffic statistics section of the output from the monitor interface t1-*fpc/pic/port* command. [PR/299688: This issue has been resolved.]
- On M320 and T-series routing platforms, when member links of a Multilink Frame Relay bundle go down and come back up, an FPC in which a Link Services Queuing (LSQ) PIC is installed might stop forwarding traffic and need to be rebooted. As a workaround, install the PICs with the member links and the LSQ PIC in the same FPC. [PR/300331: This issue has been resolved.]
- If both the key and ttl statements are included at the [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel] hierarchy level for a GRE tunnel, the TTL value might be decremented incorrectly. This can cause the ping command to fail. [PR/300956: This issue has been resolved.]
- When you configure an unnumbered interface to borrow from a loopback or non-Ethernet interface and also configure unrestricted proxy ARP on the unnumbered interface, the incoming proxy-ARP requests are dropped. As a workaround, configure the unnumbered interface to borrow from any Ethernet interface. [PR/301101: This issue has been resolved.]
- If an interface is configured as a shared uplink for the JCS 1200 platform (the shared-uplink statement is included at the [edit interfaces *interface-name*] hierarchy level), it cannot function as a label-switched interface (LSI). [PR/305520: This issue has been resolved.]
- When you enable or disable MPLS on an interface configured as a shared uplink to the JCS 1200 platform, an FPC that has a tunnel PIC installed might generate an error. [PR/305670: This issue has been resolved.]
- VPLS flood forwarding might not work correctly on an interface configured as a shared uplink to the JCS 1200 platform (the shared-uplink statement is included at the [edit interfaces *interface-name*] hierarchy level). [PR/307213: This issue has been resolved.]
- On the TX Matrix platform, if there are a large number of interface configuration changes in a small amount of time, or if the alarm process (alarmd) restarts, it might take a long time for the show version detail command to return all of its output. [PR/307228: This issue has been resolved.]
- During graceful Routing Engine switchover (GRES), resynchronization between Routing Engines might fail. In this case, the Kernel database field in the output of the show system switchover command reports the value Connection error, Initialize error. [PR/307501: This issue has been resolved.]
- When a PE router receives a PIM Join message from a CE router and the source for the required multicast data is another directly connected CE router, the attempt to create a flood next hop might initially fail. Messages including the

following are written to the system log: "NH: Failed to install flood nexthop: *index*." The next hop is eventually installed, so there is no operational impact. [PR/307579: This issue has been resolved.]

- On T-series routing platforms with VPLS configured, if a customer edge-facing interface on a provider edge router is on an Enhanced Scaling FPC4, the following message might be written repeatedly to the system log: "LCHIP(0): *number* new errors in SLout OP". The condition that triggers the message has no operational impact. [PR/309044: This issue has been resolved.]
- On M120 routers or M320 platforms with M320 Enhanced III FPCs, packets might be discarded after a graceful Routing Engine switchover event. The following might be written to the system log: "ichip_f_check_dest_errors: Fabric request time out for plane *index* dest *index* pfe *index*." To restore forwarding performance, restart the Enhanced III FPC on M320 routers or the Forwarding Engine Board on M120 routers. [PR/310061: This issue has been resolved.]
- In a Protected System Domain, under the following conditions VPLS traffic received on a core-facing shared uplink interface is not forwarded: (a) both the main routing instance and a logical system are using the shared uplink interface and (b) an FPC housing a tunnel PIC goes down and comes back up. As a workaround, configure another logical system for the main routing instance, so that all the shared uplink interfaces and peer tunnel interfaces are configured in a logical system. [PR/311302: This issue has been resolved]
- When the mirror-flash-on-disk statement is included at the [edit system] hierarchy level and the Routing Engine is rebooted, the following spurious message appears when you log in to the Routing Engine: "NOTICE: System is running on alternate media device (/dev/*device-file*)." [PR/311768: This issue has been resolved.]
- When two BGP peers are configured to use MD5 authentication and you issue the clear bgp neighbor command on one peer, the following message might be written to the system log on the other peer: "tcp_auth_ok: Packet from *address*:*identifier* missing MD5 digest." Traffic forwarding is not affected. [PR/312680: This issue has been resolved.]
- When the authentication-key statement is included at the [edit protocols bgp group *group-name*] hierarchy level, TCP sessions might not be terminated properly. As a result the message "tcp_auth_ok: Packet from *address* missing MD5 digest" might be written to the system log for each TCP ACK packet sent from a remote endpoint. [PR/313119: This issue has been resolved.]
- On MX-series and M120 routers, and M320 routers with an Enhanced III FPC, if the configuration includes the explicit-null statement at the [edit protocols mpls] or [edit protocols ldp] hierarchy level, a DPC or FPC might reboot (but not generate an error) when an MPLS packet with time-to-live (TTL) equal to 0 (zero) or 1 (one) is processed at the egress of a tunnel. [PR/313319: This issue has been resolved.]
- The output from the traceroute command includes both the IP address and DNS hostname of each hop. The hostname information might be incorrect for one or more hops. [PR/389794: This issue has been resolved.]
- During recovery after the Routing Engine hard drive fails, the JUNOS kernel might fail, causing the router to reboot. [PR/390306: This issue has been resolved.]

- When a member link of an aggregate interface goes down and comes back up and new forwarding information is installed during that change-in-status period, traffic might be lost. [PR/392550: This issue has been resolved.]
- On T-series routing platforms with aggregated SONET/SDH interfaces, if multiple statistics requests for these interfaces are queued at the same time, a memory corruption might occur, causing the kernel to crash. [PR/393572: This issue has been resolved.]

User Interface and Configuration

- Under certain conditions, when you issue show configuration | compare command the management process (mgd) might generate an error. [PR/281705: This issue has been resolved.]
- If a BGP peer is defined in a configuration group, it might not be possible to establish a connection with it. [PR/283238: This issue has been resolved.]
- If you use the replace pattern command to change the name of a policy that is applied to an object in the [edit protocols] hierarchy (for example, the import *policy-name* statement is included at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level) and then commit the configuration, the show | compare command reports the name change at the hierarchy level for the object but shows the new name as both the old and new value. The output remains the same even after multiple repetitions of the commit command. However, the policy with the new name is being applied correctly. [PR/294344: This issue has been resolved.]
- When you include the match *regular-expression* statement at the [edit system syslog (console | file | host | user)] hierarchy level to refine the set of messages included in the log, messages that do not match the expression are still included. [PR/295523: This issue has been resolved.]
- Under the following conditions, the commit operation might fail with the syntax error "inactive: group *group-name* { ... }": (a) you use the configure private command to enter configuration mode, (b) a BGP group is deactivated, and (c) you change another BGP group's name. As a workaround, use the configure command to enter configuration mode. [PR/300917: This issue has been resolved.]
- When you invoke a commit or commit check operation for a configuration that includes forwarding-table filters, the firewall process (dfwd) might generate an error and restart. [PR/301806: This issue has been resolved.]
- When TACACS+ authentication is configured and a user tries to log in to the router over an SSH or FTP connection, the JUNOS software does not include the remote user address in the authentication request packet sent to the TACACS+ server. [PR/301927: This issue has been resolved.]
- If the set of transient changes specified in a commit script (enclosed by the <transient-change> tag) includes the deactivation of a configuration statement, none of the transient changes take effect. [PR/307352: This issue has been resolved.]

Interfaces and Chassis

- On channelized T3 interfaces, the T1 loopback state does not reflect loopbacks set by facilities data link requests using the **remote-loopback-respond** statement at the [**edit interfaces** *interface-name* **t1-options**] hierarchy level. [PR/45837: This issue has been resolved.]
- If you include the **compression-device** statement at the [**edit interfaces at-***fpc/pic/port* **unit** *logical-unit-number*] hierarchy level (that is, on an ATM interface), the JUNOS kernel might generate an error and restart. [PR/265542: This issue has been resolved.]
- On 1-port 10-Gigabit Ethernet XFP Uplink PICs and 1-port 10-Gigabit Ethernet XENPAK PICs, when the 10-Gigabit Ethernet port is disabled through the CLI, the transmit laser is shut off correctly. After this, if the XFP or XENPAK module is changed or reseated, the transmit laser is turned on, even though the port is disabled. [PR/267308: This issue has been resolved.]
- When you issue the **show interfaces diagnostics optics** command and do not specify an interface name, the output is the same as for the **show interfaces** command, instead of including optic diagnostics. [PR/285978: This issue has been resolved.]
- In JUNOS Release 9.0 and later, the **monitor interface** *interface-name* command output is missing some information. [PR/296131: This issue has been resolved.]
- The commit operation does not fail when the configuration includes the following invalid combination of statements: the address specified by the **source** or **destination** statement at the [**edit interfaces gr-***fpc/pic/port* **unit** *logical-unit-number* **tunnel**] hierarchy level is the same as the interface's own subnet address (as specified by the **address** statement at the [**edit interfaces gr-***fpc/pic/port* **unit** *logical-unit-number* **family** *family-name*] hierarchy level). [PR/299443: This issue has been resolved.]
- When a Routing Engine switchover takes place, the kernel might generate an error. [PR/301327: This issue has been resolved.]
- On a router without redundant Routing Engines (such as the M7i router), if the Routing Engine restarts, the router might stop forwarding packets. As a workaround on the M7i router, issue the **request chassis cfeb restart** command. [PR/301788: This issue has been resolved.]
- On a Gigabit Ethernet IQ2 PIC with SFPs, if a logical interface is configured for VRRP, the values in the **Traffic statistics** section of the output from the **show interfaces ge-***fpc/pic/port* **extensive** command might not be accurate. [PR/303151: This issue has been resolved.]
- If you change the MTU for a shared-uplink interface on the root system domain (RSD) (by adding the **mtu** statement at the [**edit interfaces** *interface-name*] hierarchy level or changing its value), the RSD process (rsdd) generates an error and the MTU does not change. [PR/303256: This issue has been resolved.]
- In a Protected System Domain with a large number of LSPs configured (for example, 50,000), an FPC might generate an error when you issue the **show pfe route mpls** command repeatedly. [PR/303349: This issue has been resolved.]

- If you change any VRRP configuration statement (at the [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address*] hierarchy level and commit the configuration, VRRP performs a mastership election even if the changed statement does not affect mastership. [PR/303701: This issue has been resolved.]
- When you configure bandwidth management for a Protected System Domain (PSD) by including the control-plane-bandwidth-percent statement at the [edit chassis system-domains protected-system-domain psd*n*] hierarchy level, it might take up to four hours for FPC core file errors to transfer to the PSD. To reduce the transfer time to approximately 15 minutes, use one of the following workarounds: (a) remove the control-plane-bandwidth-percent statement, or (b) set the control-plane-bandwidth-percent value to 96 on the PSD to which the FPC is assigned. [PR/304765: This issue has been resolved.]
- When the links in a redundant LSQ bundle are not configured at the remote site, if a graceful Routing Engine switchover occurs and then a primary or secondary LSQ PIC goes offline, the backup Routing Engine might generate an error. [PR/306667: This issue has been resolved.]
- For SONET/SDH interfaces, when the hold-time statement is included at the [edit interfaces so-*fpc/pic/port*] hierarchy level and you change the framing type from the default (SONET) to SDH by including the framing sdh statement at the [edit interfaces so-*fpc/pic/port*] hierarchy level, the interface does not come up after the commit operation. As a workaround, deactivate the hold-time statement before changing the framing. [PR/306687: This issue has been resolved.]
- When you disable a Fast Ethernet interface, a router at the other end of a link to the interface might not mark the link as down. [PR/307538: This issue has been resolved.]
- The 1-port ATM2 OC48/STM12 IQ PIC might generate an RDI-P error when it receives a packet in which the bits corresponding to the enhanced path-RDI encoding of the G1 path overhead byte are set, even if the formal path-RDI bit within the G1 path overhead byte is not set. [PR/309929: This issue has been resolved.]
- When you set a nondefault payload size for a SAToP pseudowire (by including the payload *bytes* statement at the [edit interfaces *interface-name* satop-options] hierarchy level), the setting does not take effect and the default payload size is retained. The payload size is reported in the TDM payload size field in the output of the show route table l2circuit detail command. [PR/311066: This issue has been resolved.]
- When you configure a shared uplink interface on the JCS 1200 platform, the interface process (dcd) might generate an error and stop operating. [PR/311384: This issue has been resolved.]

Services Applications

- If Network Address Port Translation (NAPT) is configured and multiple short-lived flows are established, ports on MS PICs might not be assigned correctly. In some cases, this situation causes the MS PIC to stop functioning. [PR/300553: This issue has been resolved.]

- If Network Address Port Translation (NAPT) is configured and multiple short-lived flows are established, ports on MS PICs might not be assigned correctly. In some cases, this situation causes the MS PIC to stop functioning. [PR/304088: This issue has been resolved.]
- When a PPP session on a dedicated interface is terminated, associated static routes might remain in the routing table. [PR/309771: This issue has been resolved.]

Subscriber Access Management

- The router's address-assignment pool support enables you to create a named address range that is based on a specific DHCP option 82 value (either circuit-id or remote-id). However, when a client request is received, the router ignores the specified option 82 value and instead uses the first named range of addresses in the address-assignment pool. [PR/263077: This issue has been resolved.]
- When you configure either AAA or local authentication for Mobile IP services (at the [edit services mobile-ip] hierarchy level), a call-setup rate of more than 20 calls per second might cause the following: (a) a significant drop in the connection rate and (b) a high CPU utilization rate for the Mobile IP process (mipd) when there are more than 30,000 configured subscribers. [PR/307121: This issue has been resolved.]
- On a router configured for Mobile IP services, under the following conditions both the Mobile IP process (mipd) and the authentication process (authd) might generate an error and restart: (a) the order aaa statement is included at the [edit services mobile-ip authenticate] hierarchy level, (b) the call setup rate is more than 20 calls per second, and (c) more than 30,000 subscribers are configured. [PR/308707: This issue has been resolved.]
- On a router configured for Mobile IP services, when 40,000 concurrent subscribers are logged in, the authentication process (authd) might create an error and restart. [PR/309778: This issue has been resolved.]
- When you change a dynamic profile in the [edit dynamic-profiles] configuration hierarchy and commit the configuration, the foreign file propagation process (ffp) might generate an error. As workaround, remove the dynamic profile, commit the configuration, reinsert the dynamic profile with the desired changes, and commit again. [PR/310327: This issue has been resolved.]

Layer 2 Ethernet Services

- When you configure bridge options for a trunk interface (by including the interface statement at the [edit bridge-domain *domain-name* bridge-options] hierarchy level) and the bridge domain is part of the default virtual switch, the JUNOS software rejects the configuration as invalid. As a workaround, include the complete bridge domain configuration at the [edit routing-instances *routing-instance-name*] hierarchy level, along with another interface statement at that level for the trunk interface. [PR/307000: This issue has been resolved.]
- When you change the values for the vlan-id and vlan-tags statements at the [edit routing-instances *routing-instance-name* bridge-domains *domain-name*] hierarchy level, the multicast snooping process (mcsnoopd) might generate an error. There

is no operational effect and the process recovers automatically. [PR/307322: This issue has been resolved.]

- On an MX-series router with a large-scale Layer 2 Control Protocol configuration, Layer 2 traffic might be discarded after an in-service software upgrade. [PR/311893: This issue has been resolved.]
- On MX-series routers, access ports configured for VSTP (the interface *interface-name* statement corresponding to the port is included at the [edit protocols vstp] hierarchy level) might not interoperate properly with other vendors' switches. [PR/390026: This issue has been resolved.]

Routing Protocols

- You can specify a value for the lsp-interval statement at the [edit protocols isis *interface-name*] hierarchy level that exceeds the documented maximum (the operation does not fail when you commit such a configuration). However, values that exceed the maximum can cause unexpected behavior. [PR/41613: This issue has been resolved.]
- If the configuration includes VPNs and nonstop active routing is enabled, the following message is written repeatedly to the system log: "Error creating dynamic logical interface from sub-unit 0: No such file or directory." [PR/277005: This issue has been resolved.]
- When an IPv6 duplicate address is detected, the interface stops forwarding but IS-IS and OSPFv3 continue to announce the interface as a valid route. However, the address is unreachable and all traffic destined to or through the interface is dropped. [PR/296740: This issue has been resolved.]
- If during an LDP outage you change the value of the ldp-synchronization hold-time statement at the [edit protocols ospf area *area-id* interface *interface-name*] or deactivate the statement, OSPF might advertise the incorrect metric for the interface. [PR/303733: This issue has been resolved.]
- If during an LDP outage you change the value of the ldp-synchronization hold-time statement at the [edit protocols isis interface *interface-name*] or deactivate the statement, IS-IS might advertise the incorrect metric for the interface. [PR/304532: This issue has been resolved.]
- When you include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level, but not the graceful-restart statement at the [edit routing-options] hierarchy level, the commit operation fails with the following message: "Error in neighbor *address* of group *group-name*: graceful restart must be enabled in routing-options too." [PR/307034: This issue has been resolved.]
- On an AS boundary router or a route reflector for a VPN address family, under the following conditions VPN routes are not imported into the routing instance (VRF instance) tables: (a) the nonstop-routing statement is included at the [edit routing-options] hierarchy level, (b) routing instances are configured for locally attached VPN sites, and (c) you deactivate and reactivate the routing instance configuration. [PR/307770: This issue has been resolved.]
- When you configure a policy that causes BGP to advertise static routes that lead to unnumbered interfaces, the routing protocol process (rpd) might generate an error. [PR/308465: This issue has been resolved.]

- If a BGP notification message has an invalid value for the length of the next-hop network address field in the MP_REACH_NLRI attribute, the JUNOS software sends error code 3, subcode 1 ("Malformed Attribute List"), instead of the code specified by RFC 2858, which is code 3, subcode 9 ("Optional Attribute Error"). [PR/308628: This issue has been resolved.]
- When you re-add a previously deleted or deactivated address statement for an interface's IPv6 address on a PIM upstream neighbor (at the [edit interfaces *interface-name* unit *logical-unit-number* family inet6] hierarchy level), the addition does not register at the downstream neighbor. On the downstream neighbor, the value in the Upstream interface and Upstream neighbor fields remains unknown in the output from the show pim join extensive command. As a workaround, issue the clear pim join command. [PR/309972: This issue has been resolved.]
- If unicast routes towards a multicast source are updated via BGP static routing and an IPv6 address on a BGP peer router is deactivated and reactivated, multicast forwarding does not function correctly. [PR/386781: This issue has been resolved.]
- If the source address for IPv6 multicast traffic is resolved by a static route, information about an upstream neighbor might not be updated after a graceful Routing Engine switchover event (the value unknown appears in both the Upstream interface and Upstream neighbor fields in the output from the show pim join extensive command). [PR/389856: This issue has been resolved.]
- When a PE router receives an external LSA of type 7 (NSSA) that has a matching VPN tag or has the DN (down) bit set, it nevertheless includes the advertised route in its OSPF route calculation. According to RFC 4576, it must ignore such routes. [PR/391733: This issue has been resolved.]

MPLS Applications

- If an ingress LSP detects a routing loop (reported as Routing loop detected[*number* times] in the output from the show mpls lsp name *lsp-name* extensive command), it might stop handling traffic. [PR/293686: This issue has been resolved.]
- After some types of network events (for example, when an interface goes down and comes back up), LDP routes might be removed incorrectly from the inet.3 routing table. As a workaround, restart all LDP sessions. [PR/297144: This issue has been resolved.]
- If you include the traffic-engineering (bgp-igp-both-ribs | mpls-forwarding) statement at the [edit protocols mpls] hierarchy level for a link-protected point-to-multipoint LSP, the routing protocol process (rpd) might generate an error. [PR/303993: This issue has been resolved.]
- When a Layer 2 circuit comes back up after an interruption of network connectivity, the JUNOS software does not record the state change appropriately, and traffic is not sent through the Layer 2 circuit connection. [PR/306043: This issue has been resolved.]
- If two point-to-multipoint branch LSPs share the same incoming interface, and one of them comes up after the other during a remerge event at a transit router, the in-label for both LSPs is marked Discard, as reported by the show route table mpls.0 command. [PR/306312: This issue has been resolved.]
- When you issue the traceroute mpls ldp command, the MPLS OAM process (mplsoamd) might generate an error. [PR/307732: This issue has been resolved.]

- If an IP address is configured as both a direct LDP neighbor and a targeted LDP neighbor, and an LDP session with the neighbor repeatedly goes down and comes up again, the routing protocol process (rpd) might generate an error and stop operating. [PR/308178: This issue has been resolved.]
- If there is a single hop to an LDP neighbor and the source address of the received LDP Link Hello address is the same as the LDP Targeted Hello source address, when the LDP link neighbor and target LDP neighbor go down and come back up in a certain sequence, the Layer 2 circuit connection might remain inactive (reported as VC-Dn in the St field of the entry for the neighbor in the output from the show l2circuit connections command). To return the connection to the active state, issue the clear ldp neighbor *address* command. [PR/312672: This issue has been resolved.]

VPNs

- When a logical tunnel (lt-) interface forwards a multicast packet, it incorrectly sets the destination MAC address. [PR/304516: This issue has been resolved.]
- Including both the interface and neighbor statements in a VPLS mesh group (that is, at the [edit routing-instances *routing-instance-name* protocols vpls mesh-group *group-name*] hierarchy level) is not a valid configuration, but the commit operation does not fail. The mesh groups are not established correctly, however, as indicated in the output from the show vpls connections extensive command. [PR/304952: This issue has been resolved.]
- A dynamic change to the provider tunnel type might cause the routing protocol process (rpd) to generate an error. [PR/305081: This issue has been resolved.]
- In rare cases, changes to the encapsulation or MAC address on a PE router's CE-facing interface, followed by a nonstop active routing (NSR) event, might disrupt Layer 2 circuit communications. The show l2circuit connections command reports an MTU Mismatch (MM) status for the Layer 2 circuit connection on the remote PE router. To restore communications, on the local PE router deactivate and reactivate the l2circuit configuration stanza at the [edit protocols] hierarchy level. To avoid the error, include the ignore-mtu-mismatch statement at the [edit protocols l2circuit local-switching interface *interface-name*] hierarchy level for every interface. [PR/306453: This issue has been resolved.]

High Availability

- On a router with BGP and nonstop active routing (NSR) enabled, after a few graceful Routing Engine switchover events (for example, three or four) the routing protocols process (rpd) might generate an error and stop operating. [PR/288783]
- Following a unified in-service software upgrade (ISSU), logical tunnel interfaces might not work properly. Problems might include failure of the ping command and formation of Layer 2 forwarding loops. As a workaround, deactivate and activate the affected interfaces after the upgrade finishes. [PR/294284: This issue has been resolved.]
- During an in-service software upgrade on a TX Matrix platform, firewall counters are reset to zero (as reported by the show firewall command) at two points: when the backup Routing Engines on the routing nodes are upgraded and when FRUs

are upgraded on a newly rebooted routing node. After the second reset to zero, the counters no longer increment. [PR/305450: This issue has been resolved.]

Class of Service

- When you remove a CoS scheduler map from an interface (by removing the scheduler-map statement at the [edit class-of-service interfaces *interface-name*] hierarchy level), corresponding data structures might not be removed from Packet Forwarding Engine memory. An attempt to configure a different scheduler map on the interface might fail, as indicated by the following message in the system log: "mqchip_red_profile() no profile space available." [PR/292223: This issue has been resolved.]
- If the configured shaping rate for an interface is low (the value of the shaping-rate statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level is less than 5m), queue transmission rates do not match the configured values. [PR/305209: This issue has been resolved.]

Forwarding and Sampling

- When you include the route-accounting statement at the [edit forwarding-options family inet6] hierarchy level, the sampling process (sampled) might generate an error. [PR/291455: This issue has been resolved.]
- Under some circumstances, when you add a prefix at the [edit policy-options prefix-list *list-name*] hierarchy level, the commit operation might fail with one of the following error messages: "Check-out failed for Firewall daemon (/usr/sbin/dfwd) without details" or "configuration check-out failed." [PR/305510: This issue has been resolved.]
- When you configure Routing Engine-based sampling (by including the sampling statement at the [edit forwarding-options] hierarchy level), 4-byte AS numbers might be incorrectly reported as 2-byte numbers in the output from the monitor start sampled command. [PR/310276: This issue has been resolved.]
- If a prefix list specified at the [edit firewall family inet6 filter *filter-name* term *term-name* from source-prefix-list] hierarchy level includes an IPv4 address, the commit operation fails with the following message: "Invalid inet6 addr: '*ipv4-address*/*prefix-length*'." [PR/310299: This issue has been resolved.]
- Specifying peer as the value for the autonomous-system-type statement at the [edit forwarding-options sampling output cflowd *hostname*] hierarchy level has no effect (the exported information is the same as when the value origin is specified). [PR/310313: This issue has been resolved.]

Network Management

- When some PIC types are taken offline and brought back online, an SNMP linkUp trap is not generated for some of the logical interfaces. [PR/294667: This issue has been resolved.]
- The JUNOS software does not generate an SNMP linkDown trap when an interface's state (represented by the ifOperStatus object) changes from up to lowerLayerDown. The trap is required by RFC 2863. [PR/297829: This issue has been resolved.]

- When you issue the monitor traffic interface or tcpdump command for a logical interface on a T1 or T3 interface, the command might fail and return the following error message: "BIOCSETIF: < interface-name > : Device not configured." [PR/310814: This issue has been resolved.]
- When you enable firewall counters for IPv4 and IPv6 traffic on an interface (by including the count statement at the [edit firewall family (inet | inet6) filter *filter-name* term *term-name* then] hierarchy level and the filter *filter-name* statement at the [edit interfaces *interface-name* unit *logical-unit-number* (inet | inet6)] hierarchy level), the show snmp mib walk jnxFWCounterByteCount command might not display all of the counters. [PR/313194: This issue has been resolved.]

Related Topics

- Features in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 5
- Changes in Default Behavior and Syntax in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 31
- Errata and Changes in Documentation for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 87
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 90

Errata and Changes in Documentation for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms

Changes to the JUNOS Documentation Set

The new *JUNOS MX-series Layer 2 Configuration Guide* provides an overview of the Layer 2 functions of the MX-series routers, including configuring bridging domains, MAC address and VLAN learning and forwarding, and spanning-tree protocols. It also details the routing instance types used by Layer 2 applications. All of this material was formerly covered in the *Junos OS Routing Protocols Configuration Guide*.

Errata

This section lists outstanding issues with the documentation.

User Interface and Configuration

- The show system statistics bridge command displays system statistics on MX-series routers. [*System Basics Command Reference*]

Interfaces and Chassis

- The version of the *Junos OS Network Interfaces Configuration Guide* accessible at the JUNOS documentation home page (http://www.juniper.net/techpubs/software/junos/) corrects some errors in the version

included on the JUNOS 9.3 Documentation DVD. We recommend always accessing the online version to obtain the most current information. [*Network Interfaces*]

- In the *Junos OS Network Interfaces Configuration Guide*, Chapter 44 "Configuring IEEE 802.1ag OAM Connectivity-Fault Management", "Configuring a CFM Interface Down Action Profile Action" section states the following: “NOTE: The action profile is supported only on the physical interface level, and not on the logical interface level.” This is incorrect, and has been revised in the 9.6R1 release of the same document. The note has been replaced with the following text: “The action profile is supported on the physical interface level and the logical interface level.” [*Network Interfaces*]

Services Applications

- The sample output and description of output fields for the show services pgcp statistics gateway command in the *Junos OS System Basics and Services Command Reference* do not reflect the changes made to the command’s output in JUNOS Release 9.3R2. [*System Basics Command Reference*]

Subscriber Access Management

- Some links in the HTML output for the *JUNOS Subscriber Access Configuration Guide* do not work correctly. If you encounter any broken links, the workaround is to locate the information by browsing the HTML expanded table of contents or the HTML index.

Routing Policy and Firewall Filters

- In the *JUNOS 9.3 Configuration and Diagnostic Automation Guide*, the hyperlinks to the jcs:parse-ip function in the Table of Contents and in the second section titled “Extension functions in the junos.xml File” (https://www.juniper.net/techpubs/software/junos/junos93/swconfig-automation/extension-functions-in-the-junosxsl-file_1.html) link to the wrong target (the jcs:break-lines function). To access the reference page for the jcs:parse-ip function, navigate to the preceding or subsequent page and use the [Next] or [Prev] hyperlink. [*Automation*]
- For MX-series routers only, you can configure a firewall filter to provide matching on packet loss priority (PLP) level carried in the frame for any protocol family. The match condition can specify a single value or a range of values. [*Layer 2, Policy*]

VPNs

- When you issue a ping vpls instance command, a chassis MAC address is drawn from the ingress PE router's pool of MAC addresses and used to create the VPLS ping packet. The ping packet is then forwarded to the egress PE router. When the egress PE router receives the ping packet, it learns the MAC address from the VPLS ping packet. The MAC address is added to the egress PE router's MAC table. [*System Basics Command Reference*]

- The LDP BGP VPLS interworking feature is currently supported only on MX-series and M320 routers. [*VPNs*]
- The ability to terminate multiple Layer 2 circuit pseudowires at a single VPLS mesh group is not supported. Do not configure the local-switching statement at the [edit routing-instances *routing-instance-name* protocols vpls mesh-group *group-name*] hierarchy level. You must instead configure a mesh group for each Layer 2 circuit pseudowire terminating at the router. [*VPNs*]

High Availability

- The section titled "Nonstop Active Routing Layer 2 Circuit and LDP-Based VPLN Support" in the *Junos OS High Availability Configuration Guide* uses the term *VPLN* instead of the intended term *VPLS*. [*High Availability*]
- The ability to terminate multiple Layer 2 circuit pseudowires at a single VPLS mesh group is not supported. Do not include the local-switching statement at the [edit routing-instances *routing-instance-name* protocols vpls mesh-group *group-name*] hierarchy level. You must instead configure a mesh group for each Layer 2 circuit pseudowire terminating at the router. [*VPNs*]

Class of Service

- The *Junos OS Class of Service Configuration Guide* incorrectly states that high-priority queues on IQ2 PICs and EQ DPCs can starve lower priority queues if the high-priority queues are not rate limited. In fact, if a queue of any priority level is not rate limited, it can starve a queue with lower priority (for example, a medium-priority queue can starve a low-priority one). [*CoS*]

Mulitcast Applications

- The *Junos OS Class of Service Configuration Guide* incorrectly states that high-priority queues on IQ2 PICs and EQ DPCs can starve lower-priority queues if the high-priority queues are not rate limited. In fact, if a queue of any priority level is not rate limited, it can starve a queue with lower priority (for example, a medium-priority queue can starve a low-priority one). [*CoS*]
- In the *JUNOS Multicast Configuration Guide*, under the section titled, "Configuring PIM Join Load Balancing", the following has been updated: "When PIM join load balancing is enabled in a multicast VPN scenario with point-to-multipoint (P2MP) tunnels, the load balancing is achieved based on the join counts for the far-end provider edge (PE) routers, not for any intermediate P routers" to correctly state " When PIM join load balancing is enabled in a draft-rosen Layer 3 VPN scenario, the load balancing is achieved based on the join counts for the far-end provider edge (PE) routers, not for any intermediate P routers." In addition, the following note has been added to this section to make the support clear: "NOTE: PIM join load balancing is supported on Draft Rosen multicast VPNs (also referred to as Dual PIM Multicast VPNs). PIM join load balancing is not supported on multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast)." [*Multicast*]

System Logging

- The *destination-address* field is no longer valid in the system log message text for the following tags: RPD_IGMP_JOIN, RPD_IGMP_LEAVE, RPD_MLD_JOIN, and RPD_MLD_LEAVE. In JUNOS Release 9.3 and later, the character string "(null)" appears in the field instead of an actual address, as in this example: "RPD_IGMP_JOIN: Listener *ip-address* sent a join to (null) for group *multicast-address* source * on interface *interface-name* at *timestamp*." [*System Log*]
- When an interface configured for IGMP or MLD goes down, an RPD_IGMP_LEAVE or RPD_MLD_LEAVE message is no longer generated for each group and host pairing that is associated with the interface. Instead, a single message is generated, such as the following: "RPD_IGMP_ALL_SUBSCRIBERS_DELETED: All IGMP subscribers on interface *interface-name* deleted at *timestamp* because the interface is down."

  The RPD_IGMP_LEAVE and RPD_MLD_LEAVE messages are still generated when a subscriber session ends or times out. [*System Log*]

Related Topics

- Features in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 5
- Changes in Default Behavior and Syntax in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 31
- Issues in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 38
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 90

Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms

This section discusses the following topics:

Basic Procedure for Upgrading to Release 9.3

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 OS Installation and Upgrade Guide*.

NOTE: You cannot upgrade by more than three releases at a time. For example, if your routing platform is running JUNOS Release 8.5, you can upgrade to JUNOS Release 9.2, but not to JUNOS Release 9.3. As a workaround, first upgrade to JUNOS Release 9.2 and then upgrade to JUNOS Release 9.3.

NOTE: If both PIM and NSR are enabled on the router, you might need to perform additional steps during the upgrade to JUNOS Release 9.3. For more information, see "Upgrading to Release 9.3 on a Router Enabled for Both PIM and NSR" on page 93.

NOTE: For JUNOS Release 9.0 and later, the compact flash disk memory requirement for JUNOS software is 1 GB. For M7i and M10i routing platforms with only 256 MB memory, see the Customer Support Center JTAC Technical Bulletin PSN-2007-10-001 at https://www.juniper.net/alerts/viewalert.jsp?txtAlertNumber=PSN-2007-10-001&actionBtn=Search.

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 OS System Basics Configuration Guide*.

The download and installation process for JUNOS Release 9.3R4.4 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-9.3R4.4-domestic-signed.tgz
```

All other customers use the following command:

```
user@host> request system software add validate reboot
source/jinstall-9.3R4.4-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 9.3 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.

Upgrade Policy for JUNOS Software Extended End-Of-Life Releases

An expanded upgrade and downgrade path is now available for the JUNOS Software Extended End-of-Life (EEOL) releases. You can upgrade directly from one EEOL release to one of two adjacent later EEOL releases. You can also downgrade directly from one EEOL release to one of two adjacent earlier EEOL releases.

For example, JUNOS Software Releases 8.5, 9.3, 10.0, and 10.4 are all EEOL releases. You can upgrade from JUNOS Software Release 8.5 directly to either 9.3 or 10.0. To upgrade from Release 8.5 to 10.4, you first need to upgrade to JUNOS Software release 9.3 or 10.0, and then upgrade a second time to 10.4. Similarly, you can downgrade directly from JUNOS Software Release 10.4 to either 10.0 or 9.3. To downgrade from release 10.4 to 8.5, you first need to downgrade to 10.0 or 9.3, and then perform a second downgrade to Release 8.5.

For upgrades and downgrades to or from a non-EEOL release, the current policy is that you can upgrade and downgrade by no more than three releases at a time. This policy remains unchanged.

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

Upgrading to Release 9.3 on a Router Enabled for Both PIM and NSR

JUNOS Release 9.3 introduces NSR support for PIM for IPv4 traffic. However, the following PIM features are not compatible with NSR in this release. The commit operation fails if the configuration includes both NSR and one or more of these features:

- Anycast RP
- Draft-Rosen multicast VPNs (MVPNs)
- Local RP
- Next-generation MVPNs with PIM provider tunnels
- PIM join load balancing

JUNOS Release 9.3 introduces a new configuration statement that disables NSR for PIM only, so that you can activate incompatible PIM features and continue to use

NSR for the other protocols on the router: the nonstop-routing disable statement at the [edit protocols pim] hierarchy level. (Note that this statement disables NSR for all PIM features, not only incompatible features.)

If neither NSR nor PIM is enabled on the router to be upgraded, or only one of NSR or an incompatible PIM feature is enabled, no additional steps are necessary. Use the standard reboot or ISSU procedures described in the other sections of these instructions.

Because the new nonstop-routing disable statement is not available in JUNOS Release 9.2 and earlier, if both NSR and and an incompatible PIM feature are enabled on a router to be upgraded to JUNOS Release 9.3, you must disable either NSR or PIM before the upgrade and reenable it after the router is running JUNOS Release 9.3.

To disable and reenable NSR:

1. On the router running JUNOS Release 9.2 or earlier, enter configuration mode and disable NSR:

```
[edit]
user@host# deactivate routing-options nonstop-routing
user@host# commit
```

2. Upgrade to the JUNOS Release 9.3 software using the instructions appropriate for the router type. Note that you cannot use ISSU because it depends on NSR, which is currently disabled.
3. After the router reboots and is running JUNOS Release 9.3, enter configuration mode, disable PIM NSR, and reenable NSR:

```
[edit]
user@host# set protocols pim nonstop-routing disable
user@host# activate routing-options nonstop-routing
user@host# commit
```

To disable and reenable PIM:

1. On the router running JUNOS Release 9.2 or earlier, enter configuration mode and disable PIM:

```
[edit]
user@host# deactivate protocols pim
user@host# commit
```

2. Upgrade to the JUNOS Release 9.3 software using the instructions appropriate for the router type. You can either use the standard procedure with reboot or use ISSU.
3. After the router reboots and is running JUNOS Release 9.3, enter configuration mode, disable PIM NSR, and reenable PIM:

```
[edit]
user@host# set protocols pim nonstop-routing disable
user@host# activate protocols pim
user@host# commit
```

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a JUNOS software installation on each Routing Engine separately to avoid disrupting network operation as follows:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new JUNOS software release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the *Junos OS Installation and Upgrade Guide*.

Upgrading to Release 9.3 in a Routing Matrix

By default, when you upgrade software on the TX Matrix platform, the new image is loaded onto the TX Matrix platform and distributed to all routing nodes in the routing matrix. To upgrade software for the entire routing matrix, issue the request system software add command. Customers in the United States and Canada use the following command:

```
user@host> request system software add source/jinstall-9.3R4.4-domestic-signed.tgz
```

All other customers use the following command:

```
user@host> request system software add source/jinstall-9.3R4.4-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 TX Matrix platform.
- 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)

When you complete the software installation and reboot the TX Matrix platform, all routing nodes also reboot and all hardware and software components in the routing matrix begin using the new software.

To upgrade the backup Routing Engines, log in to the backup Routing Engine on the TX Matrix platform before you issue the request system software add command. You can also update the software on the TX Matrix platform only or on a specific T640 routing node as needed by including the lcc or scc option.

NOTE: We recommend you run the same JUNOS software release on the master and backup Routing Engines on all components of a routing matrix. If you elect to run different JUNOS software releases on the Routing Engines, a change in Routing Engine mastership can cause one or all routing nodes to be logically disconnected from the TX Matrix platform. It is also a best practice to make sure that all master Routing Engines are re0 and all backup Routing Engines are re1 (or vice versa).

NOTE: You must use the same Routing Engine model on all routing platforms in a routing matrix. For example, it is not supported to use model RE-A-2000 on the TX Matrix platform and model RE-1600 on the routing nodes.

Upgrading Using ISSU

Unified in-service software upgrade (ISSU) enables you to upgrade between two different JUNOS software releases with no disruption on the control plane and with minimal disruption of traffic. Unified in-service software upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the *Junos OS High Availability Configuration Guide*.

Downgrade from Release 9.3

To downgrade from Release 9.3 to another supported release, follow the procedure for upgrading, but replace the 9.3 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 9.3, you can downgrade the software to Release 9.0 directly, but not to Release 8.5 or earlier; as a workaround, you can first downgrade to Release 9.0 and then downgrade to Release 8.5.

For more information, see the *Junos OS Installation and Upgrade Guide*.

Related Topics

- Features in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 5
- Changes in Default Behavior and Syntax in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 31
- Issues in JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 38
- Errata and Changes in Documentation for JUNOS Software Release 9.3 for M-series, MX-series, and T-series Routing Platforms on page 87

JUNOS Software Release Notes for SRX-series Services Gateways

New Features in JUNOS Software Release 9.3 for SRX-series Services Gateways

New Features in This Release

Security

- **IKE and IPsec VPN**—JUNOS software supports a full Internet Key Exchange (IKE) and IP Security (IPsec) virtual private network (VPN) implementation. The IKE provides tunnel management for IPsec, authenticates end-entities, and performs a Diffie-Hellman key exchange to generate a VPN tunnel between network devices. The VPN tunnels generated by IKE are used to encrypt, decrypt, and authenticate user traffic between the network devices at the IP layer. In SRX-series devices, a VPN tunnel is created by distributing the IKE and IPsec workload among the multiple Services Processing Units (SPUs) of the platform. The IKE workload is distributed based on a key generated from the IKE packet's four tuples (source IP address,destination IP addresses, and UDP ports). In IPsec, the workload is distributed by the same algorithm that distributes the IKE. The Phase 2 security association (SA) for a given VPN tunnel termination points pair is exclusively owned by a particular SPU, and all IPsec packets belonging to this Phase 2 SA are forwarded to the anchoring SPU of that SA for IPsec processing.

  JUNOS software also supports the following features:

  - Site-to-site manual and IKE-negotiated (AutoKey) IPsec VPNs for policy-based VPNs and route-based VPNs
  - Authentication by preshared key and RSA public key certificates
  - Perfect Forward Secrecy (PFS) to use Diffie-Hellman Groups 1, 2, and 5
  - IKE dead peer detection (DPD) as defined in RFC 3706

- Next Hop Tunnel Binding (NHTB)—Binding multiple IPsec security associations to the same tunnel interface (This applies to static routes and OSPF only.)
- IPsec remote access with extended authentication (XAuth) support of the NS-Remote client v8.8
- VPN monitoring
- Priority queuing of IKE packets
- Invalid security parameter index (SPI) response to invalid packets
- Don't Fragment (DF) bit, including a clear option

To configure IPsec VPN options, use the ipsec statement at the [set security] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

- **Static NAT**—Static Network Address Translation (NAT) defines a one-to-one static mapping from one IP subnet to another IP subnet. To configure static NAT, use the static nat statement at the [edit security nat] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

Intrusion Detection and Prevention (IDP)

- **IDP SSL Inspection**—Secure Sockets Layer (SSL) is a protocol suite that consists of different versions, ciphers, and key exchange methods. SSLv3 and TLS protocols are supported. Combined with the Application Identification feature, the SSL Inspection feature enables SRX-series devices to inspect HTTP traffic encrypted in SSL on any TCP port. SSL inspection is disabled by default and can be enabled by using the configuration CLI. To display all installed keys and associated servers, use the show security idp ssl-inspection key command. For more information, see the *Junos OS Security Configuration Guide*.
- **IDP custom attacks and groups**—JUNOS CLI support is available for creating IDP custom attacks and groups. In previous 9.2Rx releases, creating Signature, Anomaly, and Chain custom attacks and groups required modifying XML strings. Now you can use the JUNOS configuration statements to configure the required fields.

J-Web

- **J-Web Infrastructure**—This release of JUNOS software includes revisions to the J-Web graphical user interface. The following layout and navigational elements have changed:

  The changes affect the following:

  - Dashboard
  - Menu layout
  - Configuration pages

- Monitoring pages
- Maintenance pages
- Troubleshooting sections
- Wizards

For more information, see the *Junos OS Administration Guide for Security Devices*.

Management and Administration

SNMP

- JUNOS software for SRX-series devices supports the Simple Network Management Protocol (SNMP), which is a part of the Internet protocol suite that is used to monitor network-attached devices for conditions that warrant administrative attention.

JUNOS for SRX-series Services Gateways Product Overview

Hardware

This release of JUNOS software supports the SRX 5600 and SRX 5800 services gateways, which are high-performance, highly scalable, carrier-class devices featuring multiprocessor architecture optimized for JUNOS software.

By installing different combinations of Input/Output Cards (IOCs) and Services Processing Cards (SPCs), you can tailor both the number of Gigabit ports and the maximum security processing capacity to suit your network.

The following table compares the SRX 5600 and SRX 5800 services gateways:

	SRX 5600 Device	SRX 5800 Device
Maximum Throughput	60 Gigabits per second	120 Gigabits per second
Total Slots	8	14
Slots for SPCs and IOCs	6	12
Slots for Switch Control Boards (SCBs)	2	3
Chassis Height	8 U (14”)	16 U (28”)
Devices per Rack	6	3

Two types of IOCs are available, both of which consist of four Packet Forwarding Engines and enable a throughput of 10 Gbps:

- A 40-port Gigabit Ethernet IOC with SFP connectors (1000 Mbit copper and fiber only)
- A 4-port 10-Gigabit Ethernet IOC with XFP connectors

The SRX 5600 services gateway chassis provides redundancy and resiliency. The hardware system is fully redundant, including power supplies, fan trays, and Switch Control Boards (SCBs).

Flow and Processing

- **Flow-based stateful processing**—In addition to packet processing, JUNOS software for SRX-series devices performs flow-based stateful processing. When a packet enters the device, the system applies any packet-based filter processing associated with the interface to the packet. Next, the system attempts to match the packet against an existing session based on a session's match criteria (source and destination addresses, source and destination ports, and protocol and session tokens derived from the zone and virtual router). If a packet matches an existing session, the system processes it according to the flow's session features, security policies, screens, and other features. If the packet does not match an existing

session, the system establishes a new session for the packet based on routing, policy, and other classification information. Before a packet leaves the device, the system applies filters and traffic shaping to it.

- **Distributed multithread flow**—The SRX-series services gateway is multicore, multichassis hardware with distributed computing engines. The Network Processing Units (NPUs) and multicore Services Processing Units (SPUs) on the Services Processing Cards (SPCs) comprise the data plane.

  Packets for any given flow could traverse two NPUs and possibly more than one SPU (in the case of tunnels). Therefore, a distributed flow module is needed that can span multiple computing engines.

  To configure flow options, use the flow statement at the [set security] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

Interfaces and Routing

- **Interfaces**—Interfaces act as a doorway through which traffic enters and exits a device. Several security-related configuration and runtime attributes are kept in an interface object. Different modules in the data path use these attributes. Many interfaces can share exactly the same security requirements; however, different interfaces can also have different security requirements for inbound and outbound (I/O) data packets.

  Security processing and inbound and outbound (I/O) data packets analysis are separated in JUNOS software and SRX-series service gateways. As a result, the line-card interface on the Input/Output Card (IOC) and the security processors on the Services Processing Card (SPC) are separated by a fabric. The security data plane is simultaneously performing multiprocessing (32-way MT per XLR SPU) and distributed processing (The SRX 5600 and SRX 5800 devices distribute the processing over a maximum of 2 SPUs per SPC.) For more information, see the *Junos OS Interfaces and Routing Configuration Guide*.

- **Routing**— SRX-series services gateways support using the Border Gateway Protocol (BGP), the Open Shortest Path First (OSPF) Protocol, and the Routing Information Protocol (RIP) to deliver routing information across networks. To configure the services gateway to use these protocols, use the bgp, ospf, or rip statements (respectively) at the [protocols] hierarchy level. You can also configure the services gateway to use static routes. For more information, see the *Junos OS Interfaces and Routing Configuration Guide*.

  SRX-series services gateways also support the following additional routing functionality:

  - **DHCP**— JUNOS software for SRX-series devices supports Dynamic Host Configuration Protocol (DHCP) client, relay, and server functions, enabling the services gateway to provide IP addresses and settings to hosts that are connected to the device's interfaces. When you configure the SRX-series device as a DHCP server, hosts can connect to the device's interface via subnet or through DHCP relay. To configure DHCP, use the dhcp statement at the [system services] hierarchy level.

  - **NTP**— JUNOS software for SRX-series devices incorporates Network Time Protocol (NTP) support, enabling the services gateway to synchronize time

and coordinate time distribution in a large, diverse network. To configure NTP, use the ntp statement at the [system] hierarchy level.

For more information, see the *Junos OS Administration Guide for Security Devices*.

NOTE: This release of JUNOS software for the SRX-series services gateway does not support packet-based protocols such as MPLS, Connectionless Network Service (CLNS), and IP version 6 (IPV6) and Multicast.

- **IPv4**—JUNOS software for SRX-series devices supports processing IPv4 (IP version 4) traffic through an interface. The IPv4 protocol family supports 32-bit addresses and subnets. To enable the IPv4 protocol for an interface, specify inet for the interface family. For example, use edit interfaces ge-0/0/3 unit 0 family inet address 10.10.10.10/24.
- **Class of service (CoS)** —The JUNOS software for SRX-series devices class of service (CoS) feature provides a set of mechanisms that you can use to provide differentiated services when best-effort traffic delivery is insufficient. When a network experiences congestion and delay, some packets must be dropped. CoS allows you to classify and then divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to rules that you configure. Note that CoS policing is not available in this release.

  You can use an SRX-series services gateway to control traffic rate by applying classifiers and shapers. To configure CoS components, use the component you want to configure at the [edit class-of-service] hierarchy level of the configuration. For more information, see the *Junos OS Interfaces and Routing Configuration Guide*.

Chassis Clustering

- **Chassis clustering**—You can connect a pair of the same kind of supported SRX-series devices into a cluster to provide stateful failover of JUNOS processes and services. Interchassis clustering removes the single point of failure in the network by allowing the devices to be configured in a redundant cluster, with one device acting as the primary and the other as a backup. If the primary fails, the backup takes over traffic processing. Clustered devices synchronize configuration, kernel, and Packet Forwarding Engine session states across the cluster to facilitate high availability of interfaces and services. JUNOS software includes the following chassis cluster features:
  - Resilient system architecture includes a single control plane for the entire cluster to manage multiple Packet Forwarding Engines.
  - Configuration and dynamic runtime states are synchronized between the services gateways within a cluster.

- Graceful restart of the routing protocols enables the services gateway to minimize traffic disruption during a failover.
- Physical interfaces are grouped and monitored to trigger failover to the backup services gateway if the failure parameters cross a configured threshold.

For more information, see the *Junos OS Security Configuration Guide*.

NOTE: In this release of JUNOS software for SRX-series devices, synchronization of IDP-specific runtime data does not occur across the cluster. As a result, IDP processing is not continued for sessions that fail over. (IDP processing resumes for sessions created after failover.)

NOTE: When configuring chassis clusters, you are automatically in configure private mode. As a result, you must commit changes from the top of the hierarchy. For information about the configure private mode, see the *JUNOS CLI User Guide*.

Security

- **Security zones**—Security zones are the building blocks for policies; they are logical entities to which one or more interfaces are bound. Security zones provide a means of distinguishing groups of hosts (user systems and other hosts, such as servers) and their resources from one another in order to apply different security measures to them. From the perspective of security policies, traffic enters into one security zone (to-zone) and goes out on another (from-zone). To configure security zones, use the zones statement at the [security zones] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.
- **Security policies**—Security policies can be configured to control traffic flow from one zone to another by defining a certain action on the kinds of traffic that is allowed from specified sources to specified destinations at scheduled times. When packets match a policy, the policy instructs the flow to apply different rules for features. To configure a policy, use the screen statement at the [set security policies] hierarchy level.
- **Firewall screens**—JUNOS software for SRX-series devices provides various detection methods and defense mechanisms to combat the following security breaches at all stages of their execution:
  - SYN, UDP, and ICMP flood attacks
  - Network DoS attacks
  - Operating system-specific DoS attacks

  To configure screen options, use the screen statement at the [set security screen] hierarchy level.
- **Firewall user authentication** —Firewall user authentication enables you to restrict and permit access to protected resources behind a firewall based on a

user's source IP address and other credentials. You may use pass-through authentication or Web authentication to control access to the protected resources. With pass-through authentication, a user from one zone tries to access resources from another zone over an FTP, Telnet, or HTTP connection. With Web authentication, a user tries to connect to an IP address on the device over an HTTP connection. With both methods, the device forwards the user's credentials to the server of your choice (local, RADIUS, LDAP, or RSA SecurID) to authenticate the user and control subsequent access requests.

To configure pass-through authentication, use the following statements:

set security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication pass-through

To configure Web authentication, use the following statements:

set security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication web-authentication

For more information, see the *Junos OS Security Configuration Guide*.

- **Network Address Translation**—Network Address Translation (NAT) is a method by which IP addresses in a packet are mapped from one group to another and, optionally, port numbers in the packet are translated into different port numbers. NAT is described in RFC 1631 to solve IP (version 4) address depletion problems. On an SRX-series services gateway, JUNOS software decouples NAT configuration from policy configuration. NAT has its own rules to regulate traffic on the SRX-series services gateway.

  To configure NAT, use the nat statement at the [set security] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

Intrusion Detection and Prevention (IDP)

- **IDP policies**—Intrusion Detection and Prevention (IDP) policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through an IDP-enabled device. It allows you to define policy rules to match a section of traffic based on a zone, network, and application, and then take active or passive preventive actions on that traffic.

  A policy is made up of *rulebases*, and each rulebase contains a set of *rules*. You define rule parameters, such as traffic match conditions, action, and logging requirements and then add the rules to rulebases. You can create new IDP policies from scratch, or start with a predefined template provided by Juniper Networks. Juniper Networks also provides custom application objects and attack objects that you can configure as match conditions in policies.

  To configure an IDP policy, use the idp-policy statement at the [edit security idp] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

NOTE: Installing the IDP signature database requires a license.

- **IDP signature database**—Signature database is one of the major components of IDP. It contains definitions of different objects—such as attack objects, application signatures objects, and service objects—that are used in defining IDP policy rules. As a response to new vulnerabilities, Juniper Networks periodically provides a file containing attack database updates on the Juniper Website.

  To protect your network from new threats, you can download signature database updates manually or configure your device to download them automatically at a specified interval. For more information, see the *Junos OS Security Configuration Guide*.

- **IDP application identification**—Juniper Networks provides predefined application signatures that detect TCP and UDP applications running on non standard ports. Identifying these applications allows IDP to apply appropriate attack objects to applications running on non standard ports. It also improves performance by narrowing the scope of attack signatures for applications without decoders.

  Application signatures are available as part of the security package provided by Juniper Networks. You download predefined application signatures along with the security package updates. Application identification is enabled by default and is automatically turned on when you configure the default application in the IDP policy. For more information, see the *Junos OS Security Configuration Guide*.

- **IDP protocol detector engine**—The IDP protocol detector engine contains Application Layer protocol decoders or services. You can download the protocol detector updates along with the signature database updates.

  IDP supports 52 protocol decoders or services. Protocol decoders scan protocol headers and message body to identify individual fields in the protocols to determine if data conforms to the RFC. You configure protocol decoders in IDP

policy rules to specify the protocol that an attack uses to access your network. For more information, see the *Junos OS Security Configuration Guide*

- **IDP logging**—The basic JUNOS system logging continues to function after IDP is enabled. An IDP-enabled device supports basic JUNOS system logging and continues to record events that occur because of routine operations, such as a user login into the configuration database. It records failure and error conditions, such as failure to access a configuration file. In addition to the regular system log messages, IDP generates event logs for attacks. To manage attack log volume and message size, IDP supports log suppression.

  Enabling log suppression ensures that minimal numbers of logs are generated for the same event or attack that occurs multiple times. To configure log suppression, use the suppression statement at the [edit security idp sensor-configuration log] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

- **IDP DiffServ marking**—Configuring Differentiated Services Code Point (DSCP) values in IDP policies provides a method of associating class-of-service (CoS) values—thus different levels of reliability—for different types of traffic on the network. DSCP is an integer value encoded in the 6-bit field defined in IP packet headers. It is used to enforce CoS distinctions. CoS allows you to override the default packet-forwarding behavior and assign service levels to specific traffic flows.

  You can configure DSCP value as an action in an IDP policy rule. Based on the DSCP value, behavior aggregate classifiers set the forwarding class and loss priority for the traffic determining the forwarding treatment the traffic receives. For more information, see the *Junos OS Security Configuration Guide*.

- **IDP J-Web support**—You can configure IDP policies and request security package updates by using Quick Configuration pages in the J-Web user interface. You can also display IDP status and memory usage in the J-Web monitoring pages. For more information, see the *Junos OS Security Configuration Guide* and the *Junos OS Administration Guide for Security Devices*.

Application Layerl Gateways (ALGs)

- **FTP ALGs**— JUNOS software for SRX-series devices provides File Transfer Protocol (FTP) support for services and applications that transfer data using FTP, allowing legitimate FTP traffic to go through the device while blocking out malicious FTP packets. The FTP ALG monitors PORT, PASV, and 227 commands. It performs Network Address Translation (NAT) of the IP or port in the message and gate opening on the device as necessary.

  To configure FTP ALG, use the edit security alg ftp statement at the [edit security alg] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

- **TFTP ALGs**— JUNOS software for SRX-series devices provides Trivial File Transfer Protocol (TFTP) support for services and applications that transfer data using TFTP, allowing legitimate TFTP traffic to go through the device while blocking out malicious TFTP packets. The TFTP ALG processes the TFTP packets that initiate the request and opens a pinhole to allow return packets from the reverse direction to the port that sends the request.

To configure TFTP ALG, use the edit security alg tftp statement at the [edit security alg] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

J-Web

J-Web user interface—A graphical user interface enables you to configure, monitor, troubleshoot, and manage the SRX-series devices through an Internet browser. The J-Web interface includes Quick Configuration pages to perform basic configuration of the devices and monitoring tools to view system health, routes, and statistics. The J-Web interface provides diagnostic tools (such as ping and traceroute) and file utilities to manage configuration files, licenses, and temporary files on the device. The J-Web interface also includes a chassis viewer, which provides a graphical, dynamic view of the SRX-series of devices. For more information, see the *J-Web Interface User Guide*.

Management and Administration

- **Chassis management**—JUNOS software for SRX-series devices provides the ability to monitor and manage select chassis components. This includes monitoring chassis clusters, component temperature and cooling systems, chassis firmware, and chassis location. The CLI also provides commands for bringing most chassis components online and offline.

  To bring chassis components online and offline, use the chassis statement at the [request] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

NOTE: In SRX-series services gateways, the offline, online, and restart commands are supported only on IOCs and are not supported on SPCs.

- **System logging**—JUNOS software for SRX-series devices generates separate system log messages (also called syslog messages) to record events that occur on the system's data and control planes.

  The data plane logs primarily include a list of security events that the system has handled directly inside the data plane. Because the system has already handled these events, it does not send them on to the Routing Engine. Instead, the system streams the logs directly to external log servers, bypassing the Routing Engine. To view the data plane logs, use the log statement at the [security] hierarchy level.

  The control plane logs, on the other hand, include a list of actionable events. The system sends this list of control plane events on to the eventd process on the Routing Engine, which then handles the events by using JUNOS event policies and/or by generating system log messages. You can choose to send control plane logs to a file, user terminal, routing platform console, or remote machine.

  To generate control plane logs, use the syslog statement at the [system] hierarchy level. For more information, see the *Junos OS Administration Guide for Security Devices*.

NOTE: In SRX-series devices, data plane logs and control plane logs have to be configured separately.

- **Packet tracing**—The JUNOS software for SRX-series devices trace function provides a tool for applications to write security and security flow debugging information to a file. The information that appears in this file is based on configured criteria. This criteria include source port, destination port, protocol, interface, and string matching. Use this information to analyze security application issues. The trace function operates in a distributed manner, with each thread writing to its own trace buffer. These trace buffers are then collected at one point, sorted, and written to trace files. Trace messages are delivered using the InterProcess Communications (IPC) protocol.

  To configure trace options, use the traceoptions statement at the [set security] hierarchy level. For more information, see the *Junos OS Security Configuration Guide*.

- **SPU monitoring** —JUNOS software for SRX-series devices provides a new JUNOS software-based security device that uses multiple processors to process traffic. SPU monitoring allows for:

  - CPU utilization per SPU in percentage
  - Memory utilization per SPU in percentage

  These metrics provide information that can be used to prevent unexpected outages and look for trends for capacity planning. To monitor the Flexible PIC Concentrator (FPC) card by using the SPU unit's CPU and memory utilization, use the show security monitoring fpc statement.

Related Topics

- Known Limitations in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 109
- Issues in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 111
- Errata in Documentation for JUNOS Software Release 9.3 for SRX-series Services Gateways on page 114
- Unsupported CLI Statements and Commands in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 110
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for SRX-series Services Gateways on page 115

Known Limitations in JUNOS Software Release 9.3 for SRX-series Services Gateways

Intrusion Detection and Prevention (IDP)

- This release of JUNOS software for SRX-series devices supports only the following IDP policies:

- Recommended IDP policy template
- DNS_server IDP policy template
- Custom IDP policy with Critical and Major attack groups

- IDP supports up to 100 MB for the policy size. You can create a policy with different attacks or attack groups until the size of the policy reaches 100 MB.

System

By default, the detector embedded in the SRX-series devices has the SIP, SSL, SSH, and MSPRC protocol decoders disabled.

Related Topics

- New Features in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 98
- Issues in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 111
- Errata in Documentation for JUNOS Software Release 9.3 for SRX-series Services Gateways on page 114
- Unsupported CLI Statements and Commands in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 110

Unsupported CLI Statements and Commands in JUNOS Software Release 9.3 for SRX-series Services Gateways

- On an SRX-series device, the show chassis command, does not support the pem option. [PR/296111]
- On an SRX-series device, the flow statement does not support the early ageout feature. [PR/301176]
- On an SRX-series device, the policy-option statement does not support the level, multicast-scope, dvmrp, esis, isis, l2circuit, l2vpn, ldp, msdp, ospf3, pim, ripng, rsvp, lsp-next-hop, and non-lsp-next-hop options. [PR/304324]
- On an SRX-series device, the set class-of-service routing-instances abc classifiers command does not support the exp option. [PR/305353]
- On an SRX-series device, the CLI does not support ah, egp, asp, gre, pim, rtarget, tunnel, rsvp, label-switched-path, ccc, dvrmp, esis, isis, l2circuit, l2vpn, ldp, mpls, msdp, ospf3, pim, ripng, and rsvp. [PR/305862]
- On an SRX-series device, the forwarding table filter is not supported. [PR/306854]
- On SRX-series devices, the set class-of-service and show class-of-service commands do not support the fragmentation-maps option. [PR/309209]
- On SRX-series devices, the class-of-service configuration does not support ieee-802.1ad. [PR/310006]
- On SRX-series devices, the show command does not support the oam, dot1x, subscribers, link-management, and vpls options. [PR/313099]

- On SRX-series devices, the set protocols/show protocols statement does not support the dot1x, ilmi, l2iw, lacp, link-management, neighbor-discovery, route-target, l2vpn, msdp, mvpn, ospf3, ripng, vrf, no-vrf-advertise, provider-tunnel, route-distinguisher, vrf-export, vrf-import, vrf-table-label, vrf-target, fate-sharing, dynamic-tunnels, ipv6, label-switched-path, ah, egp, asp, gre, pim, andrsvp and options. Also, igmp-snooping, l2vpn, ldp, and router-discovery are not supported under the set routing-instances abc protocols. [PR/388468]
- On SRX-series devices, the class-of-service feature does not support the interface-set option. [PR/391365]
- On SRX-series devices, the l2iw and multicast-snooping-options commands are not supported. [PR/396456]
- On SRX 5600 and SRX 5800 devices, the set security authentication-key-chains command is not supported. [PR/398127]

Related Topics

- New Features in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 98
- Known Limitations in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 109
- Issues in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 111
- Errata in Documentation for JUNOS Software Release 9.3 for SRX-series Services Gateways on page 114

Issues in JUNOS Software Release 9.3 for SRX-series Services Gateways

Outstanding Issues

Authentication

- If after the user is authenticated, the webauth-policy is deleted or changed and an entry exists in the firewall authentication table, then an authentication entry created as a result of webauth will be deleted only if a traffic flow session exists for that entry. Otherwise, the webauth entry will not get deleted and will only age out. This behavior will not cause a security breach. [PR/309534]

Chassis Clustering

- Configuring an SRX-series device with set system process jsrp-service disable only on a primary node of the cluster causes the cluster to go into an incorrect state. [PR/292411]

- The device will crash if you use set system processes chassis-control disable for 4 to 5 minutes and then enable it. Do not use this command in chassis cluster mode. [PR/296022]

Firewall

- On SRX 5600 and SRX 5800 devices, the firewall filter applied on an interface to discard the packets does not display the correct action symbol in the firewall log. [PR/399457]

Flow

- On an SRX-series device, the show security flow session command currently does not display aggregate session information. Instead, it displays sessions on a per-SPU basis. [PR/264439]
- On an SRX-series device, when traffic matches a deny policy, sessions will not be created successfully. However, sessions are still consumed, and the Unicast-sessions and Sessions-in-use fields shown by the show security flow session summary command will reflect this. [PR/284299]
- Configuring the flow filter with the all flag might result in traces that are not related to the configured filter. As a workaround, use the flow trace flag basic with the command set security flow traceoptions flag. [PR/304083]

Hardware

- On SRX 5600 and SRX 5800 devices, the LEDs on the Routing Engine and PICs are not glowing in the Chassis View in J-Web. [PR/297693]

Intrusion Detection and Prevention (IDP)

- On an SRX-series device, during compilation of especially large policies, the idp-policy subsystem may not respond to management requests after creating a policy. [PR/279147]
- On SRX 5600 and SRX 5800 devices, when the software image is downgraded from 9.3R1 to 9.2, the IDP policy compilation fails, takes an indefinite time to finish, or becomes slow due to IDP policy cache. As a workaround, follow these steps:
  - Stop the idpd daemon by using the set system processes idp-policy disable command and commit the configuration.
  - Delete all policy cache files in the /var/db/idpd/db folder.
  - Log on to SRX device as root user, and use the following UNIX command: rm –f /var/db/idpd/db/dfa* /var/db/idpd/db/pcre* /var/db/idpd/db/cache.dbd .
  - Reboot the system.

  - Enable the idpd daemon by using the delete system processes idp-policy command and commit the configuration.
  - Ensure that the cache files are regenerated and are located in the /var/db/idpd/db folder. [PR/300428]

- On SRX-series devices, when multiple applications were specified under the edit security idp idp-policy *policy-name* rulebase-ips rule *rule-number* match application configuration, IDP will process only the very first application in the configuration. To avoid false negatives, configure only one application per rule in IDP policy. [PR/302304]]
- On SRX 5600 and SRX 5800 devices, the IDP status command show security idp status displays an error message when the device is processing heavy data traffic. [PR/388048]
- On SRX-series devices, when a large number of keys are added, the Packet Forwarding Engine may not read SSL server keys due to a memory allocation error. As a workaround, restart the Packet Forwarding Engine. [PR/388102]
- On SRX 5600 and SRX 5800 devices, the IDP status command show security idp status may fail when processing heavy traffic. As a result, IDP flow, session statistics, and packet statistics does not match firewall statistics. [PR/389501]
- On SRX 5600 and 5800 devices, the HTTPS sessions with higher data transaction sizes fail due to heavy CPU usage, which results in failure of new connections. [PR/390308]

J-Web

- On SRX-series devices, J-Web does not support the configuration and show commands of static NAT.[PR/396730]

Policies

- When the firewall and IDP policy both enable diffServ marking with a different DSCP value for the same traffic, the firewall DSCP value takes precedence and the traffic is marked using the firewall DSCP value. [PR/297437]

VPN

- On an SRX-series device, if the outgoing interface and route-to-peer address is in the virtual router's routing table, IKE negotiation will not be triggered and SA cannot be negotiated. [PR/288501]
- On an SRX-series device, the shared-IKE limit for IKE users is not enforced in this release. More IKE users than the configured shared-IKE limit can establish an IKE/IPsec tunnel. [PR/288551]
- On an SRX-series device, if the first two servers are down, CRL download fails from the third alternate configured URL. [PR/306514]
- On SRX-series devices, configuring multiple tunnels between the same gateways using NHTB is not supported. [PR/314558]

Resolved Issues

Virtual Private Network (VPN)

- On SRX-series devices, because Jumbo frames were not supported, packets (either pass-through or host-bound) larger than 1500 bytes were dropped. [PR/313977: This issue has been resolved.]

Related Topics

- New Features in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 98
- Known Limitations in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 109
- Errata in Documentation for JUNOS Software Release 9.3 for SRX-series Services Gateways on page 114
- Unsupported CLI Statements and Commands in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 110

Errata in Documentation for JUNOS Software Release 9.3 for SRX-series Services Gateways

This section lists outstanding issues with the documentation.

Hardware

- In Revision 1 of the *SRX 5600 Services Gateway Getting Started Guide* and *SRX 5800 Services Gateway Getting Started Guide*, the I/O cards (IOCs) are referred to as Dense Port Concentrators (DPCs). This error will be corrected in subsequent revisions.

MIB Support

- The JUNOS 9.3 documentation does not mention support for the enterprise-specific SPU Monitoring MIB, which was added for monitoring SRX 5600 and 5800 devices. For more information, see a downloadable version of this MIB at http://www.juniper.net/techpubs/software/junos/junos93/swconfig-net-mgmt/mib-jnx-js-spu-monitoring.txt.

Screens

- The following guides contain incorrect screen configuration instructions:
  - *Junos OS Security Configuration Guide*, “Attack Detection and Prevention” chapter
  - *Junos OS Design and Implementation Guide*, “Implementing Firewall Deployments for Branch Offices” chapter

Examples throughout both of these guides describe how to configure screen options using the set security screen *screen-name* CLI statements. Instead, you should use the set security screen ids-option *screen-name* CLI statements. All screen configuration options are located at the set security screen ids-option *screen-name*] level of the configuration hierarchy.

Related Topics

- New Features in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 98
- Known Limitations in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 109
- Issues in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 111
- Unsupported CLI Statements and Commands in JUNOS Software Release 9.3 for SRX-series Services Gateways on page 110

Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for SRX-series Services Gateways

This section discusses the following topic:

- Upgrade Policy for JUNOS Software Extended End-Of-Life Releases on page 115

Upgrade Policy for JUNOS Software Extended End-Of-Life Releases

An expanded upgrade and downgrade path is now available for the JUNOS Software Extended End-of-Life (EEOL) releases. You can upgrade directly from one EEOL release to one of two adjacent later EEOL releases. You can also downgrade directly from one EEOL release to one of two adjacent earlier EEOL releases.

For example, JUNOS Software Releases 8.5, 9.3, 10.0, and 10.4 are all EEOL releases. You can upgrade from JUNOS Software Release 8.5 directly to either 9.3 or 10.0. To upgrade from Release 8.5 to 10.4, you first need to upgrade to JUNOS Software release 9.3 or 10.0, and then upgrade a second time to 10.4. Similarly, you can downgrade directly from JUNOS Software Release 10.4 to either 10.0 or 9.3. To downgrade from release 10.4 to 8.5, you first need to downgrade to 10.0 or 9.3, and then perform a second downgrade to Release 8.5.

For upgrades and downgrades to or from a non-EEOL release, the current policy is that you can upgrade and downgrade by no more than three releases at a time. This policy remains unchanged.

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

JUNOS Software Release Notes for J-series Services Routers

Features in JUNOS Software Release 9.3 for J-series Services Routers

J-series Services Router Features

For J-series Services Routers, no new features are added for the 9.3R4 release. For information on existing features, see the following manuals:

- *J2320, J2350, J4350, and J6350 Services Router Getting Started Guide*
- *J2300, J4300, and J6300 Services Router Getting Started Guide*
- *J-series Services Router Basic LAN and WAN Access Configuration Guide*
- *J-series Services Router Advanced WAN Access Configuration Guide*
- *J-series Services Router Administration Guide*

For more information about the JUNOS Internet software that runs on Services Routers, see the manuals listed in Table 11 on page 160.

Related Topics

- Issues in JUNOS Software Release 9.3 for J-series Services Routers on page 116
- Hardware Information for JUNOS Software Release 9.3 for J-series Services Routers on page 118
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for J-series Services Routers on page 120

Issues in JUNOS Software Release 9.3 for J-series Services Routers

The following problems currently exist in J-series Services Routers. The identifier following the description is the tracking number in the Juniper Networks bug database.

Interfaces and Chassis

- On channelized E1 interfaces, you might be able to configure clocking on ds-*pim*/0/*port*:*n* interfaces, where *n* is not unit 0. This is an invalid configuration and might cause a clocking selection problem on the other channels. [PR/24722]
- For ISDN dialer interfaces, when you configure the no-keepalives statement at the [edit interfaces dl0 unit *logical-unit-number*] hierarchy level and you issue the

show interfaces dl0 command, the Link flags field might still show keepalives. [PR/58520]

- If you disable a services interface by including the disable statement at the [edit interfaces sp-*pim*/0/*port*] hierarchy level and then delete the disable statement from the configuration, IPsec service is not reset correctly. As a workaround, either issue the deactivate services command followed by the activate services command, or issue the request chassis pic offline fpc-slot *pim-slot* pic-slot 0 command followed by the request chassis pic online fpc-slot *pim-slot* pic-slot 0 command. [PR/58522]
- 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, if you configure the minimum-links statement at the [edit interfaces dl0 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 dl0 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 dl0 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]
- J4350 and J6350 Services Routers might not have the requisite data buffers needed to meet expected delay-bandwidth requirements. Lack of data buffers might degrade CoS performance with smaller-sized (500 bytes or less) packets. [PR/73054]
- On J4350 and J6350 Services Routers, when an Avaya VoIP TGM550 module is in reset state, the Services Router might not respond to show chassis commands for up to 5 seconds. [PR/78695]
- 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]
- On serial interfaces transmitting either 64-byte or 128-byte packets, the effective bandwidth falls when the interface is highly oversubscribed. [PR/235753]

Platform and Infrastructure

- 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 probes might not work. [PR/75927]
- 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,

use the U3 Launchpad Installer Tool accessible at http://www.sandisk.com/Retail/Default.aspx?CatID=1411). [PR/102645]

- On J2320, J2350, J4350, and J6350 Services Routers, when you press the F10 key to save and exit from BIOS configuration mode, the operation might not work as expected. As a workaround, use the **Save and Exit** option from the **Exit** menu. 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]
- On J2320, J2350, J4350, and J6350 Services Routers, the Clear NVRAM option in the BIOS configuration mode does not work as expected. 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 mitigate this issue, note any changes you make to the BIOS configuration so that you can revert to the default BIOS configuration as needed. [PR/237722]

Services Applications

- 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 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]

Related Topics

- Features in JUNOS Software Release 9.3 for J-series Services Routers on page 116
- Hardware Information for JUNOS Software Release 9.3 for J-series Services Routers on page 118
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for J-series Services Routers on page 120

Hardware Information for JUNOS Software Release 9.3 for J-series Services Routers

Power and Heat Dissipation Requirements for J-series PIMs

On J-series Services Routers, the system monitors the PIMs and verifies that the PIMs fall within the power and heat dissipation capacity of the chassis. If power management is enabled and the capacity is exceeded, the system prevents one or more of the PIMs from becoming active.

CAUTION: Disabling power management can result in hardware damage if you overload the chassis capacities.

You can also use CLI commands to choose which PIMs are disabled. For details about calculating the power and heat dissipation capacity of each PIM and troubleshooting procedures, see the *J2320, J2350, J4350, and J6350 Services Router Getting Started Guide*.

Supported Third-Party Hardware

The following third-party hardware is supported for use with J-series Services Routers.

USB Modem

We recommend using a Multi-Tech MultiModem MT5634ZBA-USB-V92 USB modem with J-series Services Routers.

Storage Devices

The USB slots on J-series Services Routers accept a USB storage device or USB storage device adapter with a compact flash disk installed, as defined in the *CompactFlash Specification* published by the CompactFlash Association. When the USB device is installed and configured, it automatically acts as a secondary boot device if the primary compact flash disk fails on startup. Depending on the size of the USB storage device, you can also configure it to receive any core file generated during a router failure. The USB device must have a storage capacity of at least 256 MB.

Table 3 on page 119 lists USB and compact flash storage devices supported for use with the J-series routers.

Table 3: Supported Storage Devices on the J-series Services Routers

Manufacturer	Storage Capacity	Third-Party Part Number
SanDisk—Cruzer Mini 2.0	256 MB	SDCZ2-256-A10
SanDisk	512 MB	SDCZ3-512-A10
SanDisk	1024 MB	SDCZ7-1024-A10
Kingston	512 MB	DTI/512KR
Kingston	1024 MB	DTI/1GBKR
SanDisk—ImageMate USB 2.0 Reader/Writer for CompactFlash Type I and II	N/A	SDDR-91-A15
SanDisk CompactFlash	512 MB	SDCFB-512-455
SanDisk CompactFlash	1 GB	SDCFB-1000–A10

J-series Compact Flash and Memory Requirements

Table 4 on page 120 lists the compact flash and DRAM requirements for all J-series Services Routers.

Table 4: J-series Compact Flash and DRAM Requirements

Model	Minimum Compact Flash Required	Minimum DRAM Required	Maximum DRAM Supported
J2300	256 MB	512 MB	1 GB
J2320	256 MB	512 MB	2 GB
J2350	256 MB	512 MB	2 GB
J4300	256 MB	512 MB	1 GB
J4350	256 MB	512 MB	2 GB
J6300	256 MB	512 MB	1 GB
J6350	256 MB	1 GB	2 GB

Related Topics

- Features in JUNOS Software Release 9.3 for J-series Services Routers on page 116
- Issues in JUNOS Software Release 9.3 for J-series Services Routers on page 116
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for J-series Services Routers on page 120

Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for J-series Services Routers

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 card, 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 card:
  - 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 Card" on page 128.
- 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

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 card.

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 OS 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-cf*nnn*.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.
- cf*nnn* 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 card, 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
```

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 Webpage. 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 Website 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 121.
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 card, see "Special Instructions for J-series Routers with a 256-MB compact flash Card" on page 128.

Installing Software Upgrades with the J-Web Interface

If your router has at least a 512-MB compact flash card, 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

If your router has at least a 512-MB compact flash card, 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 123.
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 5 on page 124.
4. Click **Fetch and Install Package**. The software is activated after the router has rebooted.

Table 5: 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: ftp://*hostname/pathname/package-name* http://*hostname/pathname/package-name*
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 services gateway is automatically rebooted when the upgrade is complete.	Check the box if you want the services gateway 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 device, 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 123.
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 6 on page 124.
4. Click **Upload Package**. The software is activated after the services gateway has rebooted.

Table 6: 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 **Browse** to navigate to the location.

Table 6: Upload Package Summary *(continued)*

Field	Function	Your Action
Reboot If Required	If this box is checked the services gateway is automatically rebooted when the upgrade is complete.	Select the check box if you want the services gateway 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 device, see "Special Instructions for J-series Routers with a 256-MB compact flash Card" on page 128.

This section contains the following topics:

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.

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 card 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:

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 Card

For upgrading from JUNOS Release 8.5 or any 9.0 version earlier than 9.0R3, first upgrade to 8.5R3 or 9.0R2 to change to the flash utilization that will enable the system to load JUNOS 9.2R2.

Upgrade Policy for JUNOS Software Extended End-Of-Life Releases

An expanded upgrade and downgrade path is now available for the JUNOS Software Extended End-of-Life (EEOL) releases. You can upgrade directly from one EEOL release to one of two adjacent later EEOL releases. You can also downgrade directly from one EEOL release to one of two adjacent earlier EEOL releases.

For example, JUNOS Software Releases 8.5, 9.3, 10.0, and 10.4 are all EEOL releases. You can upgrade from JUNOS Software Release 8.5 directly to either 9.3 or 10.0. To upgrade from Release 8.5 to 10.4, you first need to upgrade to JUNOS Software release 9.3 or 10.0, and then upgrade a second time to 10.4. Similarly, you can downgrade directly from JUNOS Software Release 10.4 to either 10.0 or 9.3. To downgrade from release 10.4 to 8.5, you first need to downgrade to 10.0 or 9.3, and then perform a second downgrade to Release 8.5.

For upgrades and downgrades to or from a non-EEOL release, the current policy is that you can upgrade and downgrade by no more than three releases at a time. This policy remains unchanged.

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

Related Topics

- Features in JUNOS Software Release 9.3 for J-series Services Routers on page 116
- Issues in JUNOS Software Release 9.3 for J-series Services Routers on page 116
- Hardware Information for JUNOS Software Release 9.3 for J-series Services Routers on page 118

JUNOS Software with Enhanced Services Release Notes for J-series Services Routers

Features in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers

JUNOS Software with Enhanced Services Features

Release 9.3R4 of JUNOS software with enhanced services includes the following features. For more information, see the manuals described in “List of Technical Publications” on page 160.

Switching

- **IGMP Snooping**—Internet Group Management Protocol (IGMP) snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member ports. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces. To configure IGMP snooping, use the [edit protocol igmp-snooping] statement. For more information, see the *Junos OS Interfaces and Routing Configuration Guide*.

JUNOS Features Not Supported for Chassis Clusters

For this release of JUNOS software with enhanced services, the following features are not supported when chassis clustering is enabled on the router:

- **Packet-based protocols.** All packet-based protocols, such as Multiprotocol Label Switching (MPLS), Connectionless Network Service (CLNS), and IP version 6 (IPv6)
- **Services interfaces functions.** Any function that depends on the configurable J-series services interfaces:
  - ls-0/0/0—Link services Multilink Point-to-Point Protocol (MLPPP), Multilink Frame Relay (MLFR), and Compressed Real-Time Transport Protocol (CRTP)
  - gr-0/0/0—Generic routing encapsulation (GRE) and tunneling
  - ip-0/0/0—IP-over-IP (IP-IP) encapsulation
  - pd-0/0/0, pe/0/0/0, and mt-0/0/0—All multicast protocols
  - lt-0/0/0—Real-time performance monitoring (RPM)
- **WXC Integrated Services Module (WXC ISM 200)**
- **Ethernet switching on some PIMs:**
  - 4-port Fast Ethernet ePIMs running in switching mode
  - 8-port and 16-port Gigabit Ethernet uPIMs running in switching mode
- **ISDN BRI**

Related Topics

- Features Not Supported for Chassis Clusters in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 136
- Changes in Default Behavior and Syntax in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Issues in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Hardware Requirements for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 137
- Upgrade and Downgrade Instructions for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 139

Changes in Default Behavior and Syntax in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers

The following current system behavior, configuration statement usage, and operational mode command usage might not yet be documented in the JUNOS software with enhanced services documentation:

For Security

- J-series Services Routers do not support the authentication order password radius or password ldap in the edit access profile *profile-name* authentication-order command. Instead, use the order radius password or ldap password.

Related Topics

- Features in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 129
- Features Not Supported for Chassis Clusters in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 136
- Issues in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Hardware Requirements for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 137
- Upgrade and Downgrade Instructions for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 139
- Errata in Documentation for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 135

Issues in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers

Outstanding Issues

Authentication

- Your attempt to log in to the router from a management device through FTP or Telnet might fail if you type your username and password in quick succession before the prompt is displayed, in some operating systems. As a workaround, type your username and password after getting the prompts. [PR/255024]

Chassis Cluster

- In a chassis cluster, the show interface terse command on the secondary Routing Engine does not display the same details as that of the primary Routing Engine. [PR/237982]
- Because the clear security alg sip call command triggers a SIP RTO to synchronize sessions in a chassis cluster, use of the command on one node with the node-id,

local, or primary option might result in a SIP call being removed from both nodes. [PR/263976]

- When a new redundancy group is added to a chassis cluster, the node with lower priority might be elected as primary when the preempt option is not enabled for the nodes in the redundancy group. [PR/265340]
- When you commit a configuration for a node belonging to a chassis cluster, all the redundancy groups might fail over to node 0. If graceful protocol restart is not configured, the failover can destabilize routing protocol adjacencies and disrupt traffic forwarding. To allow the commit operation to take place without causing a failover, we recommend that you use the set chassis cluster heartbeat-threshold 5 command on the cluster. [PR/265801]
- In a chassis cluster, a high load of SIP ALG traffic might result in some call leaks in active resource manager groups and gates on the backup router. [PR/268613]
- In a chassis cluster, J-Web does not enable you to configure the address book. We recommend that you use the command-line interface (CLI) to configure the address book. [PR/281986]

Class of Service

- J4350 and J6350 Services Routers might not have the requisite data buffers needed to meet expected delay-bandwidth requirements. Lack of data buffers might degrade CoS performance with smaller-sized (500 bytes or less) packets. [PR/73054]
- With a CoS configuration, when you try to delete all the flow sessions using the clear security flow session command, the WX application acceleration platform may fail over with heavy traffic. [PR/273843]
- In J2350 Services Routers, the CoS does not work when the data sent to the egress GE interface is more than 100 MB. [PR/281367]

Enhanced switching

- In case of traffic going through one of the ports of a LAG running LACP, any change in remote port (for example, port going down) does not change the distribution of traffic at the local switch. [PR/292136]
- When a native VLAN is removed from a port, it still accepts untagged traffic and untagged traffic is still transmitted out of it. Restarting chassisd corrects this behavior. [PR/299961]
- If the access port is tagged with the same VLAN that is configured at the port, the access port accepts tagged packets and determines the MAC. [PR/302635]

Flow

- OSPF over GRE over IPsec does not work. [PR/105279]
- In JUNOS software with enhanced services, the TTL value on the Internet control message protocol (ICMP) responses is set to 65. [PR/233844]
- Even when forwarding options are set to drop packets for the ISO protocol family, the router forms End System-to-Intermediate System (ES-IS) adjacencies and transmits packets because ES-IS packets are Layer 2 terminating packets. [PR/252957]
- OSPF over a multipoint interface connected as a hub-and-spoke network does not restart when a new path is found to the same destination. [PR/280771]

- On J-series Services Routers, outbound filters will be applied twice for host-generated IPv4 traffic. [PR/301199]
- When route aggregation is configured, the source mask length and the destination AS fields in the CFLOW record will be 0. [PR/308083]
- When RPF processes configuration changes, it prunes all the previously joined multicast flows and immediately rejoins them, causing momentary multicast traffic interruption. [PR/309904]

Infrastructure

- 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, use the U3 Launchpad Installer Tool accessible at http://www.sandisk.com/Retail/Default.aspx?CatID=1411). [PR/102645]
- If the router does not have an ARP entry for an IP address, it drops the first packet from itself to that IP address. [PR/233867]
- On J2320, J2350, J4350, and J6350 Services Routers, when you press the F10 key to save and exit from BIOS configuration mode, the operation might not work as expected. As a workaround, use the **Save and Exit** option from the **Exit** menu. 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]
- On J2320, J2350, J4350, and J6350 Services Routers, the Clear NVRAM option in the BIOS configuration mode does not work as expected. 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 mitigate this issue, note any changes you make to the BIOS configuration so that you can revert to the default BIOS configuration as needed. [PR/237722]
- If you enable security trace options, the log file might not be created in the default location at /var/log/security-trace. As a workaround, manually set the log file to the directory /var/log/security-trace. [PR/254563]

Interfaces and Chassis

- The link status of the onboard Gigabit Ethernet interfaces (ge-0/0/0 through ge-0/0/3) or the 1-port Gigabit Ethernet ePIM interface on J4350 and J6350 Services Routers fails when you configure these interfaces in loopback mode. [PR/72381]

Routing

- Asymmetric routing, such as tracing a route to a destination behind J-series routers running JUNOS software with enhanced services with Virtual Router Redundancy Protocol (VRRP), does not work. [PR/237589]

System

- The ping status of the generic routing interfaces (gr-x/y/x) connection established through ISDN simulator fails. As a workaround, deactivate and reactivate the generic routing interfaces. [PR/282588]

VPN

- The proxy-identity statement is valid for route-based VPN configuration only. Policy-based VPN does not support the proxy-identity statement. [PR/296468]

WXC Integrated Services Module

- When two J-series routers with WXC Integrated Services Modules (WXC ISM 200s) installed are configured as peers, traceroute fails if redirect-wx is configured on both peers. [PR/227958]
- JUNOS software with enhanced services does not support policy-based VPN with WXC Integrated Services Modules (WXC ISM 200s). [PR/281822]

Resolved Issues

Authentication

- During user authentication, the firewall authentication table in the output of the security firewall-authentication users command displayed multiple failures even though the network table in the output of show network-access requests statistics showed successful authentications. [PR/250780: This issue has been resolved.]

Chassis Cluster

- In a chassis cluster, CA certificate enrollment from the secondary Routing Engine did not work. [PR/278420: This issue has been resolved.]

Related Topics

- Features in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 129
- Features Not Supported for Chassis Clusters in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 136
- Changes in Default Behavior and Syntax in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Hardware Requirements for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 137

Errata in Documentation for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers

MIB Support

- The JUNOS 9.3 documentation does not mention support for the enterprise-specific SPU Monitoring MIB, which was added for monitoring SRX 5600 and SRX 5800 devices. For more information, see a downloadable version of this MIB at http://www.juniper.net/techpubs/software/junos/junos93/swconfig-net-mgmt/mib-jnx-js-spu-monitoring.txt .

Screens

- The following guides contain incorrect screen configuration instructions:
  - *Junos OS Security Configuration Guide*, "Attack Detection and Prevention" chapter
  - *Junos OS Design and Implementation Guide*, "Implementing Firewall Deployments for Branch Offices" chapter

  Examples throughout both of these guides describe how to configure screen options using the set security screen *screen-name* CLI statements. Instead, you should use the set security screen ids-option *screen-name* CLI statements. All screen configuration options are located in the [set security screen ids-option *screen-name*] level of the configuration hierarchy.

Related Topics

Features Not Supported for Chassis Clusters in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers

For this release of JUNOS software with enhanced services, the following features are not supported when chassis clustering is enabled on the router:

- **Packet-based protocols.** All packet-based protocols, such as Multiprotocol Label Switching (MPLS), Connectionless Network Service (CLNS), and IP version 6 (IPv6)
- **Services interfaces functions.** Any function that depends on the configurable J-series services interfaces:
  - ls-0/0/0—Link services Multilink Point-to-Point Protocol (MLPPP), Multilink Frame Relay (MLFR), and Compressed Real-Time Transport Protocol (CRTP)
  - gr-0/0/0—Generic routing encapsulation (GRE) and tunneling
  - ip-0/0/0—IP-over-IP (IP-IP) encapsulation
  - pd-0/0/0, pe/0/0/0, and mt-0/0/0—All multicast protocols
  - lt-0/0/0—Real-time performance monitoring (RPM)
- **WXC Integrated Services Module (ISM 200)**
- **Ethernet switching on some PIMs:**
  - 4-port Fast Ethernet ePIMs running in switching mode
  - 8-port and 16-port Gigabit Ethernet uPIMs running in switching mode
- **ISDN BRI**

Related Topics

- Features in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 129
- Changes in Default Behavior and Syntax in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Issues in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Hardware Requirements for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 137
- Upgrade and Downgrade Instructions for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 139
- Errata in Documentation for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 135

Hardware Requirements for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers

Power and Heat Dissipation Requirements for J Series PIMs

On J-series Services Routers, the system monitors the PIMs and verifies that the PIMs fall within the power and heat dissipation capacity of the chassis. If power management is enabled and the capacity is exceeded, the system prevents one or more of the PIMs from becoming active.

CAUTION: Disabling power management can result in hardware damage if you overload the chassis capacities.

You can also use CLI commands to choose which PIMs are disabled. For details about calculating the power and heat dissipation capacity of each PIM and troubleshooting procedures, see the *J Series Services Routers Hardware Guide*.

Supported Third-Party Hardware

The following third-party hardware is supported for use with J-series Services Routers running Junos OS.

USB Modem

We recommend using a U.S. Robotics USB 56K V.92 Modem, model number USR 5637.

Storage Devices

The USB slots on J-series Services Routers accept a USB storage device or USB storage device adapter with a CompactFlash card installed, as defined in the *CompactFlash Specification* published by the CompactFlash Association. When the USB device is installed and configured, it automatically acts as a secondary boot device if the primary CompactFlash card fails on startup. Depending on the size of the USB storage device, you can also configure it to receive any core files generated during a router failure. The USB device must have a storage capacity of at least 256 MB.

Table 7 on page 137 lists the USB and CompactFlash card devices supported for use with the J-series Services Routers.

Table 7: Supported Storage Devices on the J-series Services Routers

Manufacturer	Storage Capacity	Third-Party Part Number
SanDisk—Cruzer Mini 2.0	256 MB	SDCZ2-256-A10
SanDisk	512 MB	SDCZ3-512-A10

Table 7: Supported Storage Devices on the J-series Services Routers *(continued)*

Manufacturer	Storage Capacity	Third-Party Part Number
SanDisk	1024 MB	SDCZ7-1024-A10
Kingston	512 MB	DTI/512KR
Kingston	1024 MB	DTI/1GBKR
SanDisk—ImageMate USB 2.0 Reader/Writer for CompactFlash Type I and II	N/A	SDDR-91-A15
SanDisk CompactFlash	512 MB	SDCFB-512-455
SanDisk CompactFlash	1 GB	SDCFB-1000.A10

J Series CompactFlash and Memory Requirements

Table 8 on page 138 lists the CompactFlash card and DRAM requirements for J Series Services Routers.

Table 8: J Series CompactFlash Card and DRAM Requirements

Model	Minimum CompactFlash Card Required	Minimum DRAM Required	Maximum DRAM Supported
J2320	512 MB	512 MB	1 GB
J2350	512 MB	512 MB	1 GB
J4350	512 MB	512 MB	2 GB
J6350	512 MB	1 GB	2 GB

Related Topics

- Features in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 129
- Features Not Supported for Chassis Clusters in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 136
- Changes in Default Behavior and Syntax in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Issues in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Upgrade and Downgrade Instructions for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 139
- Errata in Documentation for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 135

Upgrade and Downgrade Instructions for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers

NOTE: This information applies only to upgrading one release of JUNOS software with enhanced services to another. To upgrade from the JUNOS software to JUNOS software with enhanced services, see the *Junos OS Migration Guide*.

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 512 MB. To use the J-Web interface for an upgrade, you must have 512 MB or more.
- Before upgrading to JUNOS software with enhanced services, perform the following:
  - Upgrade to a 512-MB compact flash. For upgrading the DRAM module or compact flash, see the "Upgrading the DRAM Module or Compact Flash" section of the *Junos OS Migration Guide*. For information on formatting a new, blank compact flash card, see the "Configuring Internal Compact Flash Recovery" section of the *Junos OS Administration Guide for Security Devices*.

This section contains the following topics:

Upgrade and Downgrade Overview

Typically, you upgrade JUNOS software with enhanced services 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 and JUNOS software with enhanced services 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 signed 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.y-distribution.tgz*.

- *package-name* is the name of the package—for example, junos-jsr.
- *m.n* is the software release, with *m* representing the major release number—for example, 9.0.
- *Z* indicates the type of software release. For example, R indicates released software, and B indicates beta-level software.
- *x.y* represents the version of the major software release—for example, 1.1.
- *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 JUNOS software with enhanced services package name is junos-jsr-9.34.1-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-jsr.
- *m.n* is the software release, with *m* representing the major release number—for example, 8.5.
- *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, 1.
- 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 JUNOS software with enhanced services recovery package name is junos-jsr-8.5R1-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 J2320 or J2350 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
```

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 Webpage. 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 Website 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 software with enhanced services image for your application. For information about JUNOS software with enhanced services packages, see “Upgrade and Downgrade Overview” on page 139.
4. Download JUNOS software with enhanced services to a local host or to an internal software distribution site.

Upgrade Policy for JUNOS Software Extended End-Of-Life Releases

An expanded upgrade and downgrade path is now available for the JUNOS Software Extended End-of-Life (EEOL) releases. You can upgrade directly from one EEOL release to one of two adjacent later EEOL releases. You can also downgrade directly from one EEOL release to one of two adjacent earlier EEOL releases.

For example, JUNOS Software Releases 8.5, 9.3, 10.0, and 10.4 are all EEOL releases. You can upgrade from JUNOS Software Release 8.5 directly to either 9.3 or 10.0. To upgrade from Release 8.5 to 10.4, you first need to upgrade to JUNOS Software release 9.3 or 10.0, and then upgrade a second time to 10.4. Similarly, you can downgrade directly from JUNOS Software Release 10.4 to either 10.0 or 9.3. To

downgrade from release 10.4 to 8.5, you first need to downgrade to 10.0 or 9.3, and then perform a second downgrade to Release 8.5.

For upgrades and downgrades to or from a non-EEOL release, the current policy is that you can upgrade and downgrade by no more than three releases at a time. This policy remains unchanged.

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

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 142
- Installing Software Upgrades by Uploading Files on page 143

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. Installing software upgrades using this method copies the software image to the router.

NOTE: This procedure applies only to upgrading one release of JUNOS software with enhanced services to another. To upgrade from the JUNOS software to JUNOS software with enhanced services, see the *Junos OS Migration Guide*.

To install software upgrades from a remote server:

1. Download the software package as described in "Downloading Software Upgrades from Juniper Networks" on page 141.
2. In the J-Web interface, select **Manage > Software > Install Package**.
3. On the Install Package Quick Configuration page, enter information into the fields described in Table 9 on page 142.
4. Click **Fetch and Install Package**. The software is activated after the router reboots.

Table 9: 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 server—one of the following: ftp://*hostname/pathname/package-name* http://*hostname/pathname/package-name*

Table 9: Install Package Summary *(continued)*

Field	Function	Your Action
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 services gateway is automatically rebooted when the upgrade is complete.	Check the box if you want the services gateway 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.

NOTE: This procedure applies only to upgrading one release of JUNOS software with enhanced services to another. To upgrade from the JUNOS software to JUNOS software with enhanced services, see the *Junos OS Migration Guide*.

To install software upgrades by uploading files:

1. Download the software package as described in "Downloading Software Upgrades from Juniper Networks" on page 141.
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 10 on page 143.
4. Click **Upload Package**. The software is activated after the router has rebooted.

Table 10: Upload Package Summary

Field	Function	Your Action
File to Upload (required)	Specifies the location of the software package on the local system.	Type the location of the software package, or click **Browse** to navigate to the location.
Reboot If Required	If this box is checked the services gateway is automatically rebooted when the upgrade is complete.	Select the check box if you want the services gateway to reboot automatically when the upgrade is complete.

Installing Software Upgrades with the CLI

This section contains the following topics:

Installing Software Upgrades by Downloading Files

To install software upgrades by downloading files to the router:

1. Download the JUNOS software with enhanced services 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.

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 with enhanced services 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:

NOTE: Juniper Networks supports direct software downgrades for a maximum of three releases.

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.

NOTE: This procedure applies only to downgrading one release of JUNOS software with enhanced services to another. To downgrade JUNOS software with enhanced services to the JUNOS software, see the *Junos OS Migration Guide*.

To downgrade software with the J-Web interface:

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, click **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 software image of JUNOS software with enhanced services 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 software image of JUNOS software with enhanced services labeled with the appropriate release.

NOTE: This procedure applies only to downgrading one release of JUNOS software with enhanced services to another. To downgrade JUNOS software with enhanced services to the JUNOS software, see the *Junos OS Migration Guide*.

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 software image of JUNOS software with enhanced services labeled with the appropriate release.

Related Topics

- Features in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 129
- Features Not Supported for Chassis Clusters in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 136
- Changes in Default Behavior and Syntax in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Issues in JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 131
- Hardware Requirements for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 137
- Errata in Documentation for JUNOS Software with Enhanced Services Release 9.3 for J-series Services Routers on page 135

JUNOS Software Release Notes for EX-series Switches

- New Features in JUNOS Software for EX-series Switches, Release 9.3 on page 147
- Outstanding and Resolved Issues in JUNOS Release 9.3 for EX-series Switches on page 151
- Errata in Documentation for JUNOS Software Release 9.3 for EX-series Switches on page 158
- Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for EX-series Switches on page 158

New Features in JUNOS Software for EX-series Switches, Release 9.3

New features in Release 9.3 of JUNOS software for EX-series switches are described on the following pages:

- 802.1X, Port Security, and VoIP on page 148
- Access Control and Port Security on page 148
- Bridging, VLANs, and Spanning Trees on page 148
- Class of Service (CoS) on page 149
- High Availability on page 149
- Interfaces on page 150
- Layer 3 Protocols on page 150
- Management and RMON on page 150
- Packet Filters on page 151
- PoE on page 151

802.1X, Port Security, and VoIP

- MAC RADIUS authentication—To permit nonresponsive hosts access to the LAN, you can configure MAC RADIUS authentication on the interface to which a nonresponsive host is connected. When the MAC address of a nonresponsive host appears on the interface, the switch consults the RADIUS server to check whether the MAC address is a permitted MAC address. If the MAC address of the nonresponsive host is configured as permitted on the RADIUS server, the RADIUS server informs the switch that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected.

  You also can configure MAC RADIUS authentication to automatically eliminate the normal 90-second delay needed for the switch to determine that a device is a nonresponsive host. All 802.1X packets received on that interface will be dropped.

- Server fail fallback—Server fail fallback allows you to specify how 802.1X supplicants connected to the switch are supported if the RADIUS authentication server becomes unavailable or sends an Extensible Authentication Protocol Over LAN (EAPOL) Access-Reject message. Server fail fallback also allows you to specify that a supplicant be moved to a specified VLAN if the switch receives an EAPOL Accept-Reject message.

Access Control and Port Security

- DHCP option 82—The switch inserts DHCP option 82 information in Layer 2 and Layer 3 packets to provide information to the DHCP server about a DHCP client's network. Option 82 suboptions are circuit ID, remote ID, and vendor ID. The feature helps to protect the switch against spoofing (forging) of IP addresses and MAC addresses and against DHCP IP address starvation. If the server or clients connect to the switch through a routed VLAN interface (RVI), the switch relays the requests to the server. If the server and clients connect to the switch on the same VLAN, the switch forwards the requests. The feature supports RFC 3046, *DHCP Relay Agent Information Option*.

- DHCP snooping—Information can now be acquired and saved in the DHCP snooping database when the switch is configured as a DHCP/BOOTP relay agent or as a DHCP server (called the "local" configuration).

Bridging, VLANs, and Spanning Trees

- Private VLANs—The private VLAN (PVLAN) features on EX-series switches allow an administrator to split a broadcast domain into multiple isolated broadcast subdomains, like a VLAN inside a VLAN. Just like regular VLANs, PVLANs are isolated on Layer 2 and require a Layer 3 device to route traffic among them.

- Q-in-Q tunneling—Q-in-Q tunneling is commonly used by service providers on Ethernet access networks to segregate customer traffic into different VLANs. In Q-in-Q , a service 802.1Q (dot1q) tag is used on the service provider network to segregate traffic into different VLANs defined by the service provider.

- Unknown unicast forwarding—Unknown unicast traffic consists of unicast packets with unknown destination MAC addresses. By default, the switch floods these unicast packets that are traveling in a VLAN to all interfaces that are members of the VLAN. Forwarding this type of traffic to interfaces on the switch can trigger a security issue. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service. This is known as a traffic storm. To prevent a storm, you can disable the flooding of unknown unicast packets to all interfaces by channeling them to a specific trunk interface.

Class of Service (CoS)

- JUNOS EZQoS—JUNOS EZQoS on EX-series switches eliminates the complexities involved in configuring class of service (CoS) across the network.
- Per-interface BA classifiers—Per-interface behavior aggregate (BA) classifiers enable you to apply different BA classifiers to each interface in the switch, which allows you to classify traffic as it enters the switch.
- Port shaping—Port shaping allows you to shape aggregate traffic through a port or channel to a rate that is less than the line or port rate. With port shaping, you can configure schedulers at the port level.
- Rate shaping—Rate shaping throttles the rate at which queues transmit packets. Rate shaping is TCP friendly; that is, it buffers packets that are above the rate, rather than dropping them.

High Availability

- MAC table aging on Virtual Chassis management VLANs—MAC table aging has been extended to the Virtual Chassis management VLAN. The aging process ensures that the switch tracks only active nodes on the network and that it can flush out nodes that are no longer available.
- Virtual Chassis fast failover—The Virtual Chassis fast failover feature is a hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link failure or a member switch failure. If a link between two members fails, traffic flow between those members must be rerouted quickly so that there is minimal traffic loss.
- Virtual Chassis software upgrade enhancements—When you upgrade software in a Virtual Chassis configuration, the upgrade will either succeed or fail on all member switches, preventing the situation in which only some Virtual Chassis member switches are upgraded.
- Virtual Chassis split and merge—If there is a disruption to the Virtual Chassis configuration due to a member switch failing or being removed from the configuration, the Virtual Chassis configuration splits into two separate Virtual Chassis. This situation could cause disruptions in the network if the two separate configurations share common resources, such as global IP addresses. The Virtual Chassis split and merge feature provides a method to prevent the split Virtual Chassis from adversely affecting the network and also allows the two parts to merge back into a single Virtual Chassis configuration after the problem that caused the split has been resolved. You can also use this feature to merge two

active but separate Virtual Chassis that have not previously been part of the same configuration into one Virtual Chassis configuration.

Interfaces

- Unicast reverse-path forwarding (RPF)—Unicast RPF helps protect the switch against denial-of-service (DoS) and distributed DoS (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface on which unicast RPF is enabled.

Layer 3 Protocols

- IPv6—Support is provided for IPv6 routing, forwarding, and management (excluding multicast).

Management and RMON

- Real-time performance monitoring—Real-time performance monitoring (RPM) enables you to configure active probes to track and monitor traffic on the switch. EX-series switches supports all JUNOS RPM options.
- sFlow technology—sFlow technology is a monitoring technology for high-speed switched or routed networks that you can use to continuously monitor traffic at wire speed on all interfaces simultaneously. sFlow data can be used to provide network traffic visibility information. JUNOS software on EX-series switches supports the sFlow standard, which is described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switches and Routed Networks*.

Packet Filters

- Additional firewall filter processing points—For Layer 2 (bridged) unicast packets, firewall filter processing points now include the egress port firewall filter and the egress VLAN firewall filter.

PoE

- Power management mode—You can use the power management mode to determine the number of interfaces that can be provided with power. The two modes of power management are static and class.

Related Topics

- Outstanding and Resolved Issues in JUNOS Release 9.3 for EX-series Switches on page 151

Outstanding and Resolved Issues in JUNOS Release 9.3 for EX-series Switches

Outstanding issues in the JUNOS Release 9.3R4 software for EX-series switches are described on the following pages. They also list the issues that have been resolved since JUNOS Release 9.2R1.

Upgrading from JUNOS Release 9.2 to Release 9.3 for EX-series Switches

Starting with JUNOS Release 9.3 for EX-series switches, during the upgrade process the switch performs reference checks on VLANs and interfaces in the 802.1X configuration stanza. If there are references in the 802.1X stanza to names or tags of VLANs that are not currently configured on the switch or to interfaces that are not configured or do not belong to the ethernet-switching family, the upgrade will fail. In addition, static MAC addresses on single-supplicant mode interfaces are not supported.

CAUTION: If your Release 9.2 configuration includes any of the following conditions, revise the configuration before upgrading to Release 9.3. If you do not take these actions, the upgrade will fail:

- Ensure that all VLAN names and tags in the 802.1X configuration stanza are configured on the switch and that all interfaces are configured on the switch and assigned to the ethernet-switching family. If the VLAN or the interface are not configured, the commit will fail.
- Remove static MAC addresses on single-supplicant mode interfaces.
- In an 802.1X configuration stanza, if authentication-profile-name does not exist and you try to commit the configuration, the commit will fail.
- In an 802.1X configuration stanza, broadcast and multicast MAC addresses are not allowed in a static MAC configuration.
- Support for static MAC address bypass in single or single-secure mode has been removed.
- In an 802.1X configuration stanza, the switch will not accept the option vrange as an assigned VLAN name.
- Enabling 802.1X and the port mirroring feature on the same interface is not supported. If you enable 802.1X and the port mirroring feature on the same interface and then attempt to commit the configuration, the commit will fail.
- In an 802.1X configuration stanza, if the VLAN name or tag specified under dot1x authenticator static does not exist and you try to commit the configuration, the commit will fail.

Downgrading from JUNOS Release 9.3 to Release 9.2 for EX 4200 Switches

When a Virtual Chassis configuration is downgraded from JUNOS Release 9.3 to Release 9.2 for EX-series switches, member switches might not retain the mastership priorities that had been configured previously. To restore the previously configured mastership priorities, commit the configuration by issuing the commit command.

Resolved Issues

Access Control and Port Security

- Occasionally, if you toggle the mode of an interface from access to trunk and then back to access, the switch might not insert the static DHCP binding entries into the DHCP snooping database. [PR/283444: This issue has been resolved.]
- In multiple supplicant mode, sometimes when you move a user who has been authenticated and logged in to a dynamic VLAN, even after the user logs off, the output of show dot1x interface shows the user as authenticated. As a workaround, issue clear dot1x interface to reset the interface. [PR/292850: This issue has been resolved.]
- When you configure a static MAC address for an interface or delete the static MAC address configuration, traffic forwarding does not occur. [PR/295898: This issue has been resolved.]
- Occasionally, DHCP snooping over multiple routed VLAN interfaces (RVIs) might fail. [PR/297479: This issue has been resolved.]

- 802.1X-authenticated clients configured using static MAC addresses are not cleared from the show dot1x interface *interface-name* output even after the link is removed. [PR/302378: This issue has been resolved.]
- After you successfully authenticate a user in multiple-supplicant mode with dynamic VLAN movement and then issue the show vlans command, the interface might not appear in the command output. [PR/304936: This issue has been resolved.]
- On EX-series switches, if you configure 802.1X and enable MAC-based VLANs and dynamic firewall filters, then after you restart 802.1X, MAC-based VLANs and dynamic firewall filters might not work as expected. [PR/305097: This issue has been resolved.]
- Beginning with JUNOS Release 9.3R2 for EX-series switches, untagged packets, BPDUs (such as in LACP and STP), and priority tagged packets are processed on logical interface 0 and not on logical interface 32767. In addition, if you have not configured any untagged interfaces, the switch creates a default logical interface 0. Logical interface 32767 no longer exists. As a result, you cannot configure a tagged interface on unit 0, and you can configure an untagged interface on unit 0 only. [PR/305338: This issue has been resolved.]
- In some cases, untrusted ports forward DHCP Request/Discover packets that contain option 82 information instead of discarding them. [PR/308678: This issue has been resolved.]
- When the switch detects an invalid MAC address on a port configured for a specific MAC address, the switch does not log this event to the system messages logs. [PR/401294: This issue has been resolved.]

Bridging, VLANs, and Spanning Trees

- When you configure VRRP on EX-series switches without specifying accept-data in the configuration and a VRRP failover occurs, traffic might be lost for about 5 minutes. As a workaround, issue the clear ethernet-switching table command on the new VRRP master. [PR/271012: This issue has been resolved.]
- When frames are switched from access to trunk interfaces (that is, when incoming frames are not tagged), the priority bits in the 802.1Q header are set to 1 by default. [PR/273079: This issue has been resolved.]
- 802.1X is not supported on private VLANs (PVLANs). [PR/294406: This issue has been resolved.]
- EX-series switches allow enabling of loop protection on interfaces that have root protection enabled. [PR/297433: This issue has been resolved.]
- Occasionally after PVLAN membership of an isolated port has been changed to a different VLAN, replicated MAC addresses are shown as being static. As a workaround, restart the Ethernet switching process (eswd). [PR/306633: This issue has been resolved.]

Class of Service

- A LAG interface configured with a custom classifier using the wildcard option is bound to a different classifier when the classifier type is applied to a single interface. [PR/293795: This issue has been resolved.]
- If a Virtual Chassis configuration splits and you try to load the factory-default configuration on the individual switches, the CoS process (cosd) might consume a large percentage of CPU activity. As a workaround, restart the CoS process. [PR/305883: This issue has been resolved.]

Infrastructure

- EX-series switches do not support interface statistics for VLAN interfaces. [PR/264501: This issue has been resolved.]
- IS-IS is not supported over routed VLAN interfaces (RVIs). [PR/269391: This issue has been resolved.]
- If you modify the configuration to change the system hostname, the name might not change when you commit the configuration. As a workaround, exit from the terminal session to the switch after you have activated the configuration, then log in again. [PR/272903: This issue has been resolved.]
- In some cases, you might not be able to disable interfaces that belong to a particular Multiple Spanning Tree Instance (MSTI). [PR/284912: This issue has been resolved.]
- Occasionally in a Virtual Chassis configuration, after a member switch becomes the master switch, you might see a license error message. If you see this error message, remove the license from the original master switch using the request system license delete *license-identifier* command. [PR/285799: This issue has been resolved.]
- The port security and 802.1X features are not supported with private VLANs (PVLANs) and Q-in-Q tunneling. [PR/299184: This issue has been resolved.]
- In some cases, loopback filters are being applied to all traffic instead of to Layer 3 traffic only. [PR/311549: This issue has been resolved.]
- Occasionally, when you create a VLAN using the vlan-range configuration statement, the unknown unicast forwarding interface cannot be created using the VLAN ID. As a workaround, do not use the vlan-range configuration statement when you create the VLAN. [PR/312364: This issue has been resolved.]
- On an EX-series switch running an automatic CoS configuration, if a switchover occurs more than two times, the switch might stop classifying packets. [PR/313538: This issue has been resolved.]
- When you issue the show interfaces command, the command output might display incorrect values for IPv6 interface statistics. [PR/396656: This issue has been resolved.]
- When you upgrade from software Release 9.0 to Release 9.3 in a Virtual Chassis configuration using a preprovisioned topology, after the upgrade you must change the pre-provisioned statement to preprovisioned in the preprovisioned configuration file. [PR/386468: This issue has been resolved.]

- If you configure a link aggregation group (LAG) interface as a native VLAN, when you restart the EX-series switch, device packets might not be forwarded. As a workaround, delete and reconfigure the LAG (aex) interface. [PR/393895: This issue has been resolved.]
- If you enable or disable a spanning-tree protocol, the switch might not generate STP-related traps. [PR/397999: This issue has been resolved.]
- The dot1qPortVlanTable object is not indexed by dot1dBasePort as stated in the MIB definition, and the values are not those of the VLAN IDs on the switch. [PR/398389: This issue has been resolved.]
- If you configure an IP address on an interface and then configure a Virtual Router Redundancy Protocol (VRRP) group whose virtual IP address is the same as the IP address of that interface, and if you then delete the VRRP group, the interface might continue to use the virtual MAC address even after the VRRP group has been deleted. [PR/398650: This issue has been resolved.]

Interfaces

- Chassis alarms do not work on the management Ethernet interface. [PR/254483: This issue has been resolved.]

Layer 2 Protocols

- IGMP snooping can process approximately only 100 IGMP leaves per second. [PR/296545: This issue has been resolved.]
- Multicast packets flood on the trunk ports until an IGMP report is received. [PR/312990: This issue has been resolved.]

Layer 3 Protocols

- In some cases, if you issue the show igmp-snooping membership detail command after a membership timeout on a port, the command output shows -1 in the Receiver count field. [PR/267781: This issue has been resolved.]
- After you issue the clear igmp-snooping static command, the invalid counter and the timeout counter might not be cleared. [PR/286495: This issue has been resolved.]
- In some cases when IGMP snooping is configured, if you change the community VLAN membership of a PVLAN, all multicast traffic is flooded. [PR/309602: This issue has been resolved.]
- When IGMP snooping has been enabled on all VLANs, multicast traffic is not flooded on a VLAN that has been enabled for Q-in-Q tunneling. [PR/393082: This issue has been resolved.]

Virtual Chassis

- When the dates on the members of a Virtual Chassis are not synchronized, a member switch or backup forwarding process (pfem) might not be able to connect to the master. [PR/278784: This issue has been resolved.]
- When two EX 4200 switches are interconnected using the two 10-gigabit uplink module ports configured as Virtual Chassis ports (VCPs) to form a Virtual Chassis, the show virtual-chassis status command output shows one of the VCPs only. This problem does not affect the functioning of the Virtual Chassis. [PR/296511: This issue has been resolved.]
- When the Virtual Chassis port (VCP) of one of the member switches in a 10-member Virtual Chassis is disabled, traffic loss occurs for more than 60 seconds. [PR/298958: This issue has been resolved.]
- Occasionally, when a Virtual Chassis configuration splits, the Virtual Chassis configuration might not merge after you reboot either some member switches or the entire Virtual Chassis. As a workaround, reboot the Virtual Chassis member switches that had split. [PR/392679: This issue has been resolved.]
- In some cases after rebooting a member switch in a Virtual Chassis configuration, the Virtual Chassis configuration might show that some of its member switches are missing. [PR/397054: This issue has been resolved.]
- In a Virtual Chassis configuration, applying a family inet firewall filter to the loopback (lo0) interface might cause communication problems between the member switches of the Virtual Chassis configuration. [PR/402722: This issue has been resolved.]

Outstanding Issues

The following issues are outstanding in the JUNOS Release 9.3R4 software for EX-series switches. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

- When you have an interface with membership in a VoIP VLAN and a guest VLAN and configured with 802.1X authentication, traffic in the VoIP VLAN is forwarded even after authentication has failed for the interface. [PR/292268]
- On EX-series switches, if you configure the RADIUS revert-interval option, the switch does not attempt to reconnect to the unreachable server after the revert interval has elapsed. [PR/304637]
- If you configure an analyzer session, for port mirroring, in which the output stanza is a VLAN, but you do not configure an input stanza, the commit will fail. As a workaround, configure an input stanza and then commit the configuration again. [PR/407559]

Bridging, VLANs, and Spanning Trees

- In a Virtual Chassis configuration with BPDU control enabled, if the Virtual Chassis undergoes a graceful Routing Engine switchover (GRES), BPDU control functionality might not work properly. [PR/285726]

Infrastructure

- The speed/duplex LED on the management Ethernet port sometimes blinks even when no cable is connected. [PR/257290]
- You cannot use the rollback rescue command to revert to a rescue configuration. As a workaround, save a known good configuration to a location from which you can reload it to your switch if needed. [PR/275480]
- If you press any key on the keyboard while the switch is rebooting, the switch enters uboot mode instead of rebooting and you see the uboot prompt (=>). If this occurs, issue the boot command at the => prompt to continue the reboot. [PR/280086]
- After you upgrade or downgrade the software on an EX-series switch (by using either the CLI or the J-Web interface), the Juniper Web Device Manager might not function properly until you clear the cache in your Web browser. [PR/286614]
- The RADIUS request sent by an EX-series switch contains both Extensible Authentication Protocol (EAP) Identity Response and State attributes. [PR/300790]
- In some cases on EX 8208 switches, OSPF, VRRP, and other control packets generated by the local CPU are not properly tagged as 802.1p packets. [PR/389276]
- If you add a VLAN as the native VLAN on a trunk port that already belongs to the same VLAN, then that port is displayed twice in the output of the show vlan *vlan-id* command. [PR/432729]
- When an Ethernet link goes down, the switch does not immediately update the alarm status. [PR/443206[
- When an EX-series switch is running the default system logging (syslog) configuration and a routed VLAN interface (RVI) with instances configured goes down and comes back up repeatedly, the switch generates unwanted debug level messages (PFE TOPO and kernel RT_PFE). [PR/465852]

Layer 2 Protocols

- In some cases, when you have a large number of VLANs and interfaces configured, a configuration change might not immediately take effect. [PR/390812]

Virtual Chassis

- In some cases, after the backup switch in a two-member Virtual Chassis configuration reboots, it assumes the role of master. As a workaround, add the set virtual-chassis no-split-detection command to the configuration. [PR/434435]

Related Topics

- New Features in JUNOS Software for EX-series Switches, Release 9.3 on page 147

Errata in Documentation for JUNOS Software Release 9.3 for EX-series Switches

This section lists outstanding issues with the documentation.

Access Control and Port Security

- When you change the MTU value that is set for an interface, the DHCP snooping database and the IP source guard database are reset.

Bridging, VLANs, and Spanning Trees

- If you configure a voice VLAN on an access interface, the interface might forward the frames tagged with an interface VLAN ID as well as packets tagged with a voice VLAN ID.

Related Topics

- New Features in JUNOS Software for EX-series Switches, Release 9.3 on page 147
- Outstanding and Resolved Issues in JUNOS Release 9.3 for EX-series Switches on page 151

Upgrade and Downgrade Instructions for JUNOS Software Release 9.3 for EX-series Switches

This section discusses the following topic:

- Upgrade Policy for JUNOS Software Extended End-Of-Life Releases on page 158

Upgrade Policy for JUNOS Software Extended End-Of-Life Releases

An expanded upgrade and downgrade path is now available for the JUNOS Software Extended End-of-Life (EEOL) releases. You can upgrade directly from one EEOL release to one of two adjacent later EEOL releases. You can also downgrade directly from one EEOL release to one of two adjacent earlier EEOL releases.

For example, JUNOS Software Releases 8.5, 9.3, 10.0, and 10.4 are all EEOL releases. You can upgrade from JUNOS Software Release 8.5 directly to either 9.3 or 10.0. To upgrade from Release 8.5 to 10.4, you first need to upgrade to JUNOS Software release 9.3 or 10.0, and then upgrade a second time to 10.4. Similarly, you can downgrade directly from JUNOS Software Release 10.4 to either 10.0 or 9.3. To downgrade from release 10.4 to 8.5, you first need to downgrade to 10.0 or 9.3, and then perform a second downgrade to Release 8.5.

For upgrades and downgrades to or from a non-EEOL release, the current policy is that you can upgrade and downgrade by no more than three releases at a time. This policy remains unchanged.

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

List of Technical Publications

Table 11 on page 160 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 12 on page 164 lists the books included in the *Network Operations Guide* series. Table 13 on page 165 lists the manuals and release notes supporting JUNOS software with enhanced services. All documents are available at http://www.juniper.net/techpubs/.

Table 14 on page 166 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 11: Technical Documentation for Supported Routing Platforms

Book	Description
JUNOS Software for Supported Routing Platforms	
Access Privilege	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.
Class of Service	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.
CLI User Guide	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 *JUNOS System Basics Configuration Guide*.
Feature Guide	Provides a detailed explanation and configuration examples for several of the most complex features in the JUNOS software.
High Availability	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).
MPLS Applications	Provides an overview of traffic engineering concepts and describes how to configure traffic engineering protocols.
Multicast Protocols	Provides an overview of multicast concepts and describes how to configure multicast routing protocols.
Multiplay Solutions	Describes how you can deploy IPTV and voice over IP (VoIP) services in your network.

Table 11: Technical Documentation for Supported Routing Platforms *(continued)*

Book	Description
MX-series Layer 2 Configuration Guide	Provides an overview of the Layer 2 functions of the MX-series routers, including configuring bridging domains, MAC address and VLAN learning and forwarding, and spanning-tree protocols. It also details the routing instance types used by Layer 2 applications. All of this material was formerly covered in the *JUNOS Routing Protocols Configuration Guide*.
MX-series Solutions Guide	Describes common configuration scenarios for the 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).
Network Interfaces	Provides an overview of the network interface functions of the JUNOS software and describes how to configure the network interfaces on the routing platform.
Network Management	Provides an overview of network management concepts and describes how to configure various network management features, such as SNMP and accounting options.
Policy Framework	Provides an overview of policy concepts and describes how to configure routing policy, firewall filters, and forwarding options.
Protected System Domain	Provides an overview of the JCS 1200 platform and the concept of Protected System Domains (PSDs). The JCS 1200 platform, which contains up to 12 Routing Engines running JUNOS software, can be connected to up to three T-series routing platforms. To configure a PSD, you assign any number of Flexible PIC concentrators (FPCs) on a 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.
Routing Protocols	Provides an overview of routing concepts and describes how to configure routing, routing instances, and unicast routing protocols.
Secure Configuration Guide for Common Criteria and JUNOS-FIPS	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.
Services Interfaces	Provides an overview of the services interfaces functions of the JUNOS software and describes how to configure the services interfaces on the router.
Software Installation and Upgrade Guide	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 *JUNOS System Basics Configuration Guide*.
Subscriber Access	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.

Table 11: Technical Documentation for Supported Routing Platforms *(continued)*

Book	Description
System Basics	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.
VPNs	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.
JUNOS References	
Hierarchy and RFC Reference	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 *JUNOS System Basics Configuration Guide*.
Interfaces Command Reference	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot interfaces.
Routing Protocols and Policies Command Reference	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot routing policies and protocols, including firewall filters.
System Basics and Services Command Reference	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.
System Log Messages Reference	Describes how to access and interpret system log messages generated by JUNOS software modules and provides a reference page for each message.
J-Web User Guide	
J-Web Interface User Guide	Describes how to use the J-Web graphical user interface (GUI) to configure, monitor, and manage Juniper Networks routing platforms.
JUNOS API and Scripting Documentation	
JUNOScript API Guide	Describes how to use the JUNOScript application programming interface (API) to monitor and configure Juniper Networks routing platforms.
JUNOS XML API Configuration Reference	Provides reference pages for the configuration tag elements in the JUNOS XML API.
JUNOS XML API Operational Reference	Provides reference pages for the operational tag elements in the JUNOS XML API.
NETCONF API Guide	Describes how to use the NETCONF API to monitor and configure Juniper Networks routing platforms.

Table 11: Technical Documentation for Supported Routing Platforms *(continued)*

Book	Description
JUNOS Configuration and Diagnostic Automation Guide	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.
Hardware Documentation	
Hardware Guide	Describes how to install, maintain, and troubleshoot routing platforms and components. Each platform has its own hardware guide.
PIC Guide	Describes the routing platform's Physical Interface Cards (PICs). Each platform has its own PIC guide.
DPC Guide	Describes the Dense Port Concentrators (DPCs) for all MX-series routers.
JUNOScope Documentation	
JUNOScope Software User Guide	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.
Advanced Insight Solutions (AIS) Documentation	
Advanced Insight Solutions Guide	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.
J-series Routing Platform Documentation	
Getting Started Guide	Provides an overview, basic instructions, and specifications for J-series Services Routers. 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 *Getting Started Guide* for your router model.
Basic LAN and WAN Access Configuration Guide	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.
Advanced WAN Access Configuration Guide	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.
Administration Guide	Shows how to manage users and operations, monitor network performance, upgrade software, and diagnose common problems on J-series Services Routers.
Release Notes	

Table 11: Technical Documentation for Supported Routing Platforms *(continued)*

Book	Description
JUNOS Release Notes	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.
Hardware Release Notes	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.
JUNOScope Release Notes	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.
AIS Release Notes	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.
AIS AI-Scripts Release Notes	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.
J-series Services Router Release Notes	Briefly describe Services Router features, identify known hardware problems, and provide upgrade and downgrade instructions.

Table 12: JUNOS Software Network Operations Guides

Book	Description
Baseline	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.
Interfaces	Describes tasks for monitoring interfaces. Tasks include using loopback testing and locating alarms.
MPLS	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.
MPLS Log Reference	Describes MPLS status and error messages that appear in the output of the show mpls lsp extensive 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.

Table 12: JUNOS Software Network Operations Guides *(continued)*

Book	Description
MPLS Fast Reroute	Describes operational information helpful in monitoring and troubleshooting an MPLS network configured with fast reroute (FRR) and load balancing.
Hardware	Describes tasks for monitoring M-series and T-series routing platforms.

To configure and operate a J-series Services Router 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 13: JUNOS Software with Enhanced Services Documentation

Book	Description
All Platforms	
JUNOS Software with Enhanced Services Interfaces and Routing Configuration Guide	Explains how to configure J-series interfaces for basic IP routing with standard routing protocols, ISDN service, firewall filters (access control lists), and class-of-service (CoS) traffic classification.
JUNOS Software with Enhanced Services Security Configuration Guide	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).
JUNOS Software with Enhanced Services Administration Guide	Shows how to monitor J-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.
JUNOS Software with Enhanced Services CLI Reference	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.
J-series Only	
JUNOS Software with Enhanced Services for J-series Services Router Design and Implementation Guide	Provides guidelines and examples for designing and implementing IPsec VPNs), firewalls, and routing on J-series Services Routers running JUNOS software with enhanced services.
JUNOS Software with Enhanced Services for J-series Services Router Quick Start	Explains how to quickly set up a J-series Services Router. This document contains router declarations of conformity.

Table 13: JUNOS Software with Enhanced Services Documentation *(continued)*

Book	Description
JUNOS Software with Enhanced Services J-series Services Router Hardware Guide	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.
JUNOS Software with Enhanced Services for J-series Services Router Migration Guide	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.
WXC Integrated Services Module Installation and Configuration Guide	Explains how to install and initially configure a WXC Integrated Services Module in a J-series Services Router for application acceleration.
JUNOS Software with Enhanced Services for J-series Services Router Release Notes	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.

Table 14: Additional Books Available Through http://www.juniper.net/books

Book	Description
Interdomain Multicast Routing	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.
JUNOS Cookbook	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.
MPLS-Enabled Applications	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.
OSPF and IS-IS: Choosing an IGP for Large-Scale Networks	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.

Table 14: Additional Books Available Through http://www.juniper.net/books *(continued)*

Book	Description
Routing Policy and Protocols for Multivendor IP Networks	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.
The Complete IS-IS Protocol	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, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/. 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 Management tool: 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 Management 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 https://www.juniper.net/cgi-bin/docbugreport/.

Revision History

03 December 2009—Revsion 6, JUNOS Release 9.3R4

28 May 2009—Revsion 5, JUNOS Release 9.3R4

13 August 2009—Revsion 4, JUNOS Release 9.3R4

15 May 2009—Revsion 3, JUNOS Release 9.3R3

18 December 2008—Revision 2, JUNOS Release 9.3R2

14 November 2008—Revision 1, JUNOS Release 9.3R1